

# Contrastive Monotonic Pixel-Level Modulation

Kun Lu<sup>✉</sup>, Rongpeng Li<sup>✉</sup>, and Honggang Zhang<sup>✉</sup>

Zhejiang University, Hangzhou, China  
{lukun199, lironpeng, honggangzhang}@zju.edu.cn

**Abstract.** Continuous one-to-many mapping is a less investigated yet important task in both low-level visions and neural image translation. In this paper, we present a new formulation called MonoPix, an unsupervised and contrastive continuous modulation model, and take a step further to enable a pixel-level spatial control which is critical but can not be properly handled previously. The key feature of this work is to model the monotonicity between controlling signals and the domain discriminator with a novel contrastive modulation framework and corresponding monotonicity constraints. We have also introduced a selective inference strategy with logarithmic approximation complexity and support fast domain adaptations. The state-of-the-art performance is validated on a variety of continuous mapping tasks, including AFHQ cat-dog and Yosemite summer-winter translation. The introduced approach also helps to provide a new solution for many low-level tasks like low-light enhancement and natural noise generation, which is beyond the long-established practice of one-to-one training and inference. Code is available at <https://github.com/lukun199/MonoPix>.

**Keywords:** Pixel-level modulation, Continuous monotonic translation, contrastive training

## 1 Introduction

Deep learning has significantly revolutionized the way we process and represent images. Though in most high-level tasks, we may safely assume that the input-ground truth pairs are deterministic and perfectly aligned [23,17,45] (e.g., class labels, bounding boxes, and key points), it does not hold true in many low-level scenarios. In low-light image enhancement [46,6] and natural noise generation [9,32] for instance, the ever-changing environment makes it even impossible to provide firm ground truth pairs. Directly adopting a one-to-one (O2O) mapping on these pseudo labels can result in several problems: in the training stage, it ignores the one-to-many (O2M) mapping property, and inherently invites certain divergence and instability; for evaluation and deployment purposes, it incurs large variance in different environments [38] and lacks certain scalability.

To enable O2M mappings, a direct solution is to introduce conditional generative models [29,7], or carry out a multi-branch ensemble on several sub-modules [47,53]. The major obstacle comes from the lack of dense and varied discrete image pairs to support the subsequent training. Some other methods approach

continuous generation by interpolating between two models, or tuning an extra adaptation module [21,62,68,22,37]. Nevertheless, methods in this category involve two correlated networks which are usually obtained from a fine-tuning or re-training. As a result, they not only prevent an end-to-end training, but also increase the storage burden and inference complexity. Moreover, all these methods can suffer from a training-inference gap that, a discrete training is not sufficient to provide a continuous modulation at inference stage.

On the other hand, through the content and style disentanglement [36,31,12], it is shown that a single generative network is capable of handling varied styles and attributes in an end-to-end manner, and provides a smooth generation by style interpolation [31,13]. To better secure the intermediate generation quality, works on continuous cross-domain translation further propose to constrain the in-between samples by coordinating two domain discriminators [18,51], or introducing some extra constraints [71]. However, to obtain a desired output, a user either has to change the random style noise at an unpredictable time complexity, or is required to provide a proper reference style which may not always be plausible. In addition, these methods are primarily focused on an image-level holistic style translation, but not typically designed for many tasks where a pixel-level modulation is preferred.

These limitations pique our interest; and we therefore devote to providing a new solution that simultaneously enables an unsupervised, end-to-end, continuous, and monotonic pixel-level image control. By injecting a spatial controlling signal into the feature space, and benefiting from the inductive bias of convolutions, we render it possible to directly manipulate from a fine-grained image-like interface. The monotonicity control is inspired by the recent progress on contrastive learning [72,10] and image-to-image translation [18,51], where we further devise a contrastive modulation scheme that connects the control monotonicity

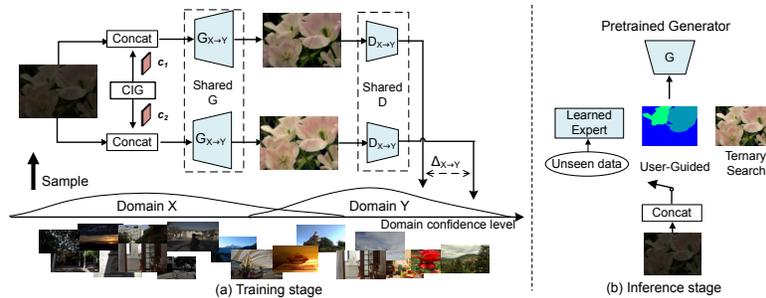


Fig. 1: Overview of MonoPix framework. Blue blocks are learnable parts. CIG denotes contrastive intensity generator (a procedure) that randomly produces two spatial control signals  $c_1$  and  $c_2$  for contrastive training.  $\Delta_{X \rightarrow Y}$  is the confidence difference. (a) During training, we model the monotonicity between control signal and domain confidence. (b) We enable multiple inference strategies or perform fast adaptation through another expert network

with the confidence level of a domain discriminator. In this way, one image more “resembles” the target domain will get rewarded with a higher discriminator score; and the training could proceed without any other style level supervisions. A similar previous work CUT [54] introduces patchwise contrastive loss (on different patches) for image translation, while it does not support image editing and O2M mapping. A schematic illustration of our approach can be found in Figure 1. Simple though these ideas are, they construct a competitive new solution for continuous domain translation, typically those in conjunction with conventional low-level tasks.

We summarize the contributions of this work:

- We investigate a new research area of continuous pixel-level image modulation which has not been particularly investigated before, and provide the corresponding solution. The surveyed topic enables an interesting unsupervised continuous translation and provides a fine-grained image editing tool.
- We propose a new contrastive learning-based solution for continuous image translation, which can be seamlessly integrated into existing GAN-based one-to-one mapping schemes, and is both training and inference-efficient. The proposed approach does not need any domain-specific labels and has no training and inference gap.
- We provide several basic evaluation metrics for continuous generation, and validate the state-of-the-art performance of the proposed model on both the translation continuity and quality. Extensive results on low-level tasks also demonstrate the effectiveness of our proposed approach.
- We are one of the first that revisit the long-established O2O mapping protocol in many low-level tasks. Our work suggests it is plausible and probably more flexible to formulate them in an O2M or M2M manner instead.

## 2 Related Work

**One-to-Many Mappings in Low-Level Task.** O2M mapping commonly occurs in low-level image processing. It is evidenced in recent literature that many tasks, such as deraining [53], noise modeling [7,21,22], and low-light image enhancement [34,64,29] can be categorized into this scope. Data-guided approaches model the diversity with synthesized one-to-many pairs (usually with physical models) and generate them with the corresponding control signals [29,7,22,5,57]. As an alternative to the dense training pairs, image modulation methods address this task in a two-stage training manner. Upchurch *et al.* [66] propose to train two networks and implement feature interpolations. In CFSNet [67] the extra network is further simplified into only a tuning branch. Inspired by the normalization operations, He *et al.* [21] propose to insert a per-layer modulation module and adjust them when trained on another enhancing level. Similarly, DynamicNet [62] performs adaptation with several extra tuning blocks. Besides feature interpolation, it is also shown that modulations can take place via interpolation on the network parameters [68] or image residuals [22]. There are also

some efforts that attempt to perform consecutive enhancement [75,58], but the increased complexity and lack of control range also limit their applications.

In contrast, we present a continuous modulation in a totally unsupervised<sup>1</sup>, end-to-end, and inference-efficient manner by training a single network that executes all these tasks. There is no training and inference gap in the proposed approach, which by design allows a continuous and realistic generation.

**Image-to-Image Translation.** CNN-based image translation mainly takes the style representation as an auxiliary loss [16,39,30,63]. Different from designing explicit style descriptors, it is observed by Dumoulin *et al.* [15] that incorporating an affine transformation after the instance normalization (IN) [65] layer can lead to a conditional style generation, which further inspires data-driven style manipulation methods like AdaIN [25] and WCT [41]. StyleBank [8], on the other hand models the style translation with a trainable bank that disentangles the style representation. It is also shown in [76] and [40] that by injecting the style information as part of the input, the model can generate varied samples with the same set of parameters.

As another line of work, GAN [19] provides an important tool for image translation. Pioneering works like Pix2Pix [27] and CycleGAN [80] introduce adversarial and unsupervised conditional generation, but are O2O mapping models. To enable a diverse translation, non-one-to-one mappings are thereafter investigated [12,48,26,3]. As one of the most popular and successful schemes, disentangling the content and style representation exhibits a great success on a wealth of attribute and style translation tasks [36,26,13,73,31], and enables a continuous translation by interpolating between two latent vectors [31,13]. Recently, works on continuous cross-domain translation further refine the quality of intermediate images by introducing an interpolation discriminator [71,43], constraining the intermediate results with discriminators from both sides [18,51], or by exploiting the path of interpolation and translation manifold [11,56,44].

Different from these existing works, the proposed approach eschews the need for either any extra discriminator or a weighted bi-lateral supervision that can hardly scale to model many-to-one mappings. Instead, we teach one discriminator by contrastively feeding two samples and modeling the consistency between control signals and translation monotonicity.

**Image Editing.** Our work is also related to image editing. This task can be achieved generally by learning a conditional image-level mapping [27,80,35,49] or by exploring the latent semantic trajectories [2,79,14,60,82]. Modern image editing methods prove to be effective in modeling multiple image-level features such as semantic contents, geometric shapes, and styles [61,42]; but there are a few that provide a spatial control [33,81,55], and typically in an unsupervised manner. The proposed method, further, provides a fine-grained, unsupervised, and spatial-level editing, whilst exhibiting both fidelity to inputs and a strong scalability for style and content manipulation.

---

<sup>1</sup> In this paper, we typically refer to methods that perform domain translation without the need for paired images and intermediate domain intensity labels.

### 3 Continuous Cross-Domain Pixel-Level Control

First we formalize a brief description of the newly examined task. Consider two domains  $X$  and  $Y$  each with corresponding image samples i.e.  $\{\mathbf{x}_i\}_{i=1}^M$  and  $\{\mathbf{y}_i\}_{i=1}^N$ , together with a control signal  $\mathbf{c}$  of the shape  $W \times H$  representing the translation intensity of each pixel. We learn a mapping function  $G_{X \rightarrow Y} : X \times C \mapsto Y$  that translate from  $X$  to  $Y$ , and from  $Y$  to  $X$  if necessary. Moreover, the translation is required to be monotonic. Thus, for each image or patch  $\mathbf{x}_i$  and two control signals  $\mathbf{c}_1$  and  $\mathbf{c}_2$ , we require the translated result  $G_{X \rightarrow Y}(\mathbf{x}_i, \mathbf{c}_2)$  better belongs to the target domain  $Y$  than its counterpart if  $c_2^{i,j} > c_1^{i,j}$  for each pixel position. In the following, we will elaborate on the detailed framework and corresponding loss functions that lead to accomplishing this goal.

#### 3.1 Pixel-Level Contrastive Generation

In existing works on unsupervised domain translation, the conditional style vector is usually constructed as a flat one [31,12] or with coarse spatial shapes [33]. These approaches are moderately satisfactory when controlling the local attributes, but can not enable a fine-grained control.

To tackle this problem, we first visit the pixel-level control signal. This idea is plain and easy to implement, where we can directly formulate the control signal as an extra input channel. Without loss of generality, we always assume that the pixel-level control signal is bounded in  $[0, 1]$ . Then, the translated result  $\hat{\mathbf{y}}$  of an input  $\mathbf{x}$  could be represented as (and vice versa from  $Y$  to  $X$ )

$$\hat{\mathbf{y}} = G_{X \rightarrow Y}(\text{concat}(\mathbf{x}, \mathbf{c})) \quad \text{for each } c^{i,j} \in [0, 1] \quad (1)$$

However, to enable such a pixel-level training is non-trivial. First, modern training schemes and datasets are only focused on holistic attributes, but there are no pixel-level labels. Second, we do not have a gallery of domain-specific intermediate labels that guide the modulation process, except for the coarse binary tags - 0 or 1, which can not model the domain distribution and precludes conventional absolute intensity-based training. To overcome these problems, we propose to formulate a contrastive learning-based approach instead, that trains the domain generator  $G$  by relative intensity. In the training stage, we first randomly generate two pixel-level control signals  $\mathbf{c}_1$  and  $\mathbf{c}_2$  for each input patch, then we encourage the discriminator to yield a higher domain-specific confidence level for that with a higher translation strength. Since the discriminator  $D$  is trained to examine both the quality and domain-specific features of generated images, a higher confidence level for  $D$  means a better likelihood that the output belongs to the target domain. This is different from DLOW [18], who needs to train two discriminators and overlooks the underlying data distributions in the source domain. Note that in the training stage, we empirically set a control signal  $\mathbf{c}$  as one filled with the same value (i.e.,  $c^{1,1} = c^{1,2} = \dots = c^{W,H}$ ), considering the randomness of control signal and the inductive bias of convolutional operation. On the contrary, assigning each  $c^{i,j}$  with randomly generated numbers during training is tedious, and also violates the smoothness of neighboring pixels.

### 3.2 Normalization-Compatible Translator

Although concatenating the pixel-level control signal  $\mathbf{c}$  can be direct and simple, it is incompatible with commonly used instance normalization modules [65]. Denoted by  $\{\mathcal{K}^i\}_{i=1}^C$  a convolution kernel where  $i$  represents the  $i$ -th channel,  $\text{Idx}$  the channel index function and  $\otimes$  the convolution operation, the convoluted result for  $\text{concat}(\mathbf{x}, \mathbf{c})$  can be calculated as  $\mathbf{x} \otimes \mathcal{K}^{i \neq \text{Idx}(\mathbf{c})} + \mathbf{c} \otimes \mathcal{K}^{i = \text{Idx}(\mathbf{c})}$ . As a result, the numerical difference between two contrastive samples is expected to be  $(\mathbf{c}_2 - \mathbf{c}_1) \otimes \mathcal{K}^{i = \text{Idx}(\mathbf{c})}$ , which becomes a constant when  $\mathbf{c}$  is element-wise the same (in convolutions, a kernel is shared across the spatial dimension). It could also be concluded that any linear operation before instance normalization would cause the same problem (for example adding the broadcast control signal to the input) as the constant difference on two channels would be eliminated by IN. To prevent this, we introduce a simple non-linear transformation like leaky ReLU [50] when adopting instance normalization. Visualizations on the activated domain-specific features (by identifying the positions where a non-linear transformation takes place) can be found in the supplement. Batch normalization, on the other hand, introduces certain contrastivity inside the training samples but does not impede the overall monotonicity.

### 3.3 Objective Functions

Based on the proposed contrastive pixel-level modulation framework, we provide the corresponding constraints steering towards a monotonic and cross-domain control. Since the translation can be bi-directional, we will always take that from  $X$  to  $Y$  as an example.

**Monotonicity Loss.** We employ the output value of target discriminator to model the likelihood of belonging to this domain. Given an input image  $\mathbf{x}$  and two pixel-level control signal  $\mathbf{c}_1$  and  $\mathbf{c}_2$  where we assume  $c_2^{i,j} > c_1^{i,j}$ , the confidence difference  $\Delta_{\mathbf{X} \rightarrow \mathbf{Y}}^{\text{tar}}$  between translated results can be calculated as:

$$\Delta_{\mathbf{X} \rightarrow \mathbf{Y}}^{\text{tar}} = D_Y(G_{X \rightarrow Y}(\mathbf{x}, \mathbf{c}_2)) - D_Y(G_{X \rightarrow Y}(\mathbf{x}, \mathbf{c}_1)) \quad (2)$$

Under the monotonicity constraint  $c_2^{i,j} > c_1^{i,j}$ , we encourage the confidence difference  $\Delta_{\mathbf{X} \rightarrow \mathbf{Y}}^{\text{tar}}$  to be positive for any  $\mathbf{c}_1$  and  $\mathbf{c}_2$ :

$$\mathcal{L}_{\text{mono-X2Y}} = \mathbb{E}_{\mathbf{x} \sim X, \mathbf{c}_1, \mathbf{c}_2 \sim C} \|\max(\epsilon - \Delta_{\mathbf{X} \rightarrow \mathbf{Y}}^{\text{tar}}, 0)\|_2^2 \quad (3)$$

where  $\epsilon > 0$  is a hyperparameter controlling the strength of monotonicity, which we call a *margin* here. As illustrated in Figure 2, the above loss exerts a quadratic punishment when  $\Delta_{\mathbf{X} \rightarrow \mathbf{Y}}^{\text{tar}} < \epsilon$ , and thus introduces not only a positive  $\Delta_{\mathbf{X} \rightarrow \mathbf{Y}}^{\text{tar}}$ , but also a positive margin. Omitting  $\epsilon$  can cause a trivial solution where the confidence difference is only a little greater than zero and finally a weak control.

**Domain Fidelity Loss.** Although monotonicity loss enables a continuous cross-domain control, there exists an asymmetry issue. As we know that a well-trained discriminator can only produce a high confidence level for a *real* image, constraining the confidence level to be low, as in the case of  $\mathbf{c}_1$ , may not be promising.

The discrepancy can lead to a biased control that lacks in the diversity and control quality. Though this can be mitigated by applying a weak paired loss (e.g.,  $L1$  or  $MSE$ ) between the input image and the translated one, a rigid identity mapping also introduces certain instability for  $\mathbf{c}_2$  and involves much parameter tuning. As a surrogate, we not only require a higher confidence level for  $\mathbf{c}_2$  from the target domain discriminator, but also a higher confidence level for  $\mathbf{c}_1$  from the source discriminator. We formulate the domain fidelity loss as:

$$\mathcal{L}_{df-X2Y} = \mathbb{E}_{\mathbf{x} \sim X, \mathbf{c}_1, \mathbf{c}_2 \sim C} \|\max(\epsilon - \Delta_{X \rightarrow Y}^{src}, 0)\|_2^2 \quad (4)$$

where  $\Delta_{X \rightarrow Y}^{src}$  is defined as:

$$\Delta_{X \rightarrow Y}^{src} = D_X(G_{X \rightarrow Y}(\mathbf{x}, \mathbf{c}_1)) - D_X(G_{X \rightarrow Y}(\mathbf{x}, \mathbf{c}_2)) \quad (5)$$

**GAN Loss.** MonoPix also utilizes the adversarial loss [19] and cycle consistency loss [80] to provide a proper gradient when training the bi-directional translation network. GAN loss is specialized in constraining the overall quality:

$$\min \max \mathcal{L}_{adv-X2Y} = \mathbb{E}_{\mathbf{y} \sim Y} [D_Y(\mathbf{y})] + \mathbb{E}_{\mathbf{x} \sim X, \mathbf{c} \sim C} [1 - D_Y(\hat{\mathbf{y}})] \quad (6)$$

We adopt cycle consistency loss since in many low-level tasks we are more concentrated on the intensity and stability of enhancement instead of the diversity of appearance. To this end, we expect that the input image could be recovered from a reverse translation with the same intensity  $\mathbf{c}$ :

$$\mathcal{L}_{cyc-X2Y} = \mathbb{E}_{\mathbf{x} \sim X} [\|G_{Y \rightarrow X}(G_{X \rightarrow Y}(\mathbf{x}, \mathbf{c}), \mathbf{c}) - \mathbf{x}\|_1] \quad (7)$$

**Overall Objective.** To keep the presentation succinct, we take the translation from  $X$  to  $Y$  as an example when constructing these loss functions. Similarly losses in the reverse direction could also be formulated. For the generator  $G = \{G_{X \rightarrow Y}, G_{Y \rightarrow X}\}$  and discriminator  $D = \{D_X, D_Y\}$ , the overall training objective is expressed as:

$$\begin{aligned} \min G \quad \mathcal{L}_G &= \mathcal{L}_{adv} + \lambda_{cyc} \mathcal{L}_{cyc} + \lambda_{mn} \mathcal{L}_{mono} + \lambda_{df} \mathcal{L}_{df} \\ \min D \quad \mathcal{L}_D &= \mathcal{L}_{adv} + \lambda_{mn} \mathcal{L}_{mono} + \lambda_{df} \mathcal{L}_{df} \end{aligned} \quad (8)$$

where  $\lambda_{cyc}$ ,  $\lambda_{mn}$  and  $\lambda_{df}$  are hyperparameters balancing different losses.

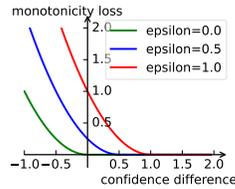


Fig. 2: Example curves of the monotonicity loss. A positive  $\epsilon$  encourages a margin on confidence difference

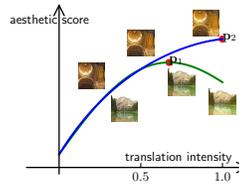


Fig. 3: Illustration of the aesthetic gap.  $\mathbf{p}$  denotes the best balance point

### 3.4 Selective Inference

MonoPix naturally generates a gallery of candidate results for a given input by constraining the monotonicity. However, a high translation intensity, and hence a high domain discriminator’s confidence level does not necessarily yield a satisfactory result. As shown in Figure 3, an image can much “resemble” the target domain but is inconsistent with human perception (aesthetic score).

We assume there lies a perfect balance point  $\mathbf{p}$  that yields the best aesthetic score, where this score can be any reasonable metric, such as human perception, and quantitative metrics like PSNR and SSIM [69] when a reference image is available. To approximate  $\mathbf{p}$ , MonoPix enables a variety of inference strategies with different execution complexity:

**Exhaustive Inference.** Benefiting from the monotonicity, we can always sample  $N$  control signals and choose the best result. The complexity is  $O(N)$ .

**Ternary Search.** Since the translation intensity - aesthetic score curve can be a monotonically increasing one or consists of two continuous monotonic parts, it is *unimodal-like*. Finding the local maxima can thus be efficiently solved by using *ternary search*. The details of this method are described in the supplementary material. It has the  $O(\log N)$  complexity.

**Guidance from Learned Expert.** We can further decrease the complexity to  $O(1)$  by introducing an expert network  $P$  that mimics target distribution. It can be trained in a few-shot manner and even on unseen datasets for fast adaptation.

## 4 Experiments and Results

We carry out experiments on both the commonly used domain translation datasets AFHQ cat-dog [13] and Yosemite summer-winter [80], and low-level tasks including prevalent LOL low-light enhancement [70] and SIDD noise generation [1]. For a fair comparison, we will always report the results when  $\mathbf{c}$  is set to be element-wise the same when not specified.

### 4.1 Metrics

For continuous cross-domain image generation, we evaluate both the dynamic translation process and the quality of translated images:

**AL/Rg:** As we expect that an image should change monotonically from the input one, the foremost concern comes from the absolute linearity (AL) and control range (Rg). For an input image, we first modulate it by varying the intensity level  $\mathbf{c}$  from 0 to 1 with an interval of 0.1. Then, AL score can be evaluated by the *Pearson correlation coefficient* between these intensity levels and the LPIPS distances [77] against the input image. Rg score is calculated by the *range* of these values, which reflects both the fidelity and difference.

**RL/Sm:** Besides absolute linearity, we follow [71] to measure the smoothness of adjacent intermediate pairs. To both overcome the drawback of *standard deviation* as used in [71] (which is sensitive to the absolute input value), and keep

in comparison with AL, we calculate relative linearity (RL) from the linearity of cumulative adjacent LPIPS distances, where a high RL score indicates the translation is even between adjacent samples. The smoothness (Sm) can be evaluated by measuring the maximum value of these adjacent LPIPS distances.

**ACC/FID:** Following [51,71], we adopt a pre-trained ResNet [23] binary classifier to measure the success of translation. For each image, we report the highest confidence level of the classifier among the consecutive modulation process. The quality of translation can be evaluated by Fréchet inception distance (FID) [24] between the whole trajectory (with interval of 0.1) and real images, as in [44].

We use the corresponding paired evaluation metrics in low-level tasks. For low-light image enhancement, we adopt PSNR, SSIM, and LPIPS. For blind noise generation, we leverage average KL divergence (AKLD) [74].

## 4.2 Implementation Details

Following [28], we use the U-net [59] structure as the generator and discriminator, and Adam as the optimizer with  $\beta_1$  and  $\beta_2$  set to be 0.5 and 0.9 respectively. LSGAN [52] is adopted to provide the basic GAN loss. As to the hyperparameter settings, we always set  $\lambda_{cyc}$  as 10, as in CycleGAN, and  $\lambda_{mn}$  as 1, which is the same as the weight of  $\mathcal{L}_{adv}$ .  $\lambda_{df}$  and  $\epsilon$ , along with training epochs vary in different datasets, which can be found in detail from the supplementary material. The learning rate is set to be  $10^{-4}$  initially and linearly drops to zero.

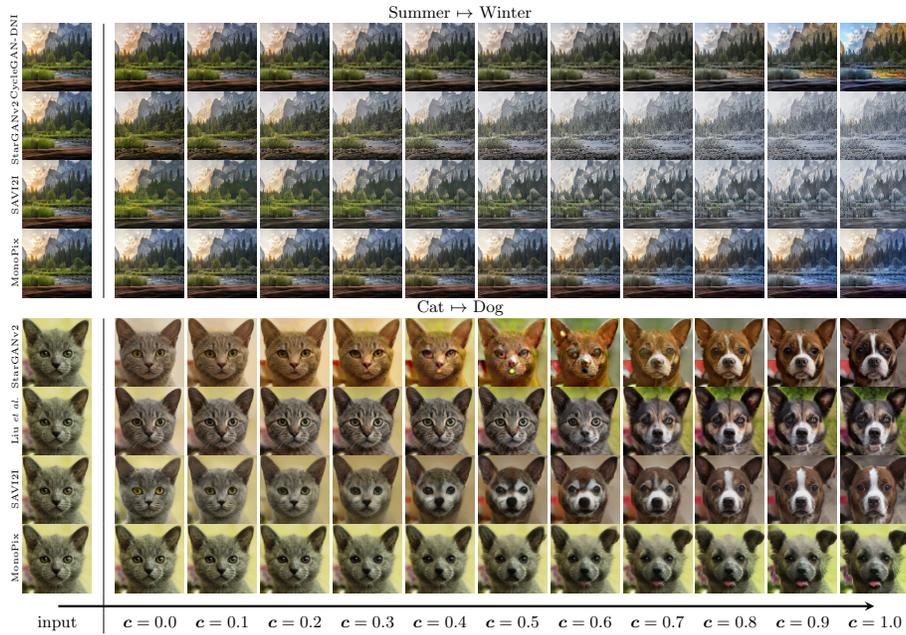


Fig. 4: Visual comparisons on continuous domain translation

### 4.3 Results on Continuous Cross-Domain Translation

We compare MonoPix with deep network interpolation method CycleGAN-DNI [68], style interpolation method StarGANv2 [13], and recently proposed continuous cross-domain translation model Liu *et al.* [44] and SAVI2I [51]. For style-guided methods, we randomly select their reference images from the target domain to provide the style information and report the average scores in 5 experiments (we call this variant “-Rdm”), and calculate the average target style information (style centroid) so that they can generate deterministic interpolations with  $O(1)$  complexity (we call this variant “-Cent”).

The visual and quantitative results can be found in Figure 4 and Table 1 respectively. We first notice that two-stage interpolation method CycleGAN-DNI and our contrastive modulation approach MonoPix show a better fidelity on the input images, while style-guided solutions bear an obvious distortion on the original textures and attributes. Though CycleGAN-DNI exhibits the highest control range in Table 1, the intermediate smoothness is significantly worse than the rest candidates. We also find that MonoPix provides a competitive absolute linearity, but the modifications between adjacent translation intensities are not strongly guaranteed to be evenly distributed, as in Yosemite scenario. This is partially related to the non-linearity of domain discriminator, but is not a key practical problem since we can execute a finer exploration in “intensity of interests”, which we illustrate more in the supplementary material. As to the realness and quality, MonoPix also performs favorably against other state of the arts, and strongly demonstrates the benefits of introducing a contrastive modulation scheme that helps to explore domain-specific features. Nevertheless, we note that evaluating FID on the whole trajectory, or on a specific intensity may produce different results for MonoPix, as we do not assume all input images share the same optimal translation intensity. More discussions on the qualitative

Table 1: Quantitative results on cross-domain continuous translation

Methods	Summer $\mapsto$ Winter			Winter $\mapsto$ Summer		
	AL $\uparrow$ /Rg $\uparrow$	RL $\uparrow$ /Sm $\downarrow$	ACC $\uparrow$ /FID $\downarrow$	AL $\uparrow$ /Rg $\uparrow$	RL $\uparrow$ /Sm $\downarrow$	ACC $\uparrow$ /FID $\downarrow$
CycleGAN-DNI [68]	0.969/0.286	0.937/0.067	0.766/35.0	0.979/0.279	0.911/0.079	0.834/42.8
StarGANv2-Rdm [13]	0.805/0.134	0.992/0.021	0.616/46.1	0.741/0.107	0.996/0.018	0.771/50.1
StarGANv2-Cent [13]	0.719/0.100	0.985/0.016	0.650/49.3	0.417/0.055	0.996/0.010	0.720/52.8
SAVI2I-Rdm [51]	0.959/0.185	0.994/0.013	0.680/44.5	0.939/0.168	0.995/0.013	0.706/47.6
SAVI2I-Cent [51]	0.934/0.145	0.988/0.010	0.718/47.5	0.873/0.111	0.993/0.009	0.721/49.2
MonoPix (Ours)	0.939/0.154	0.940/0.015	0.950/38.1	0.960/0.145	0.964/0.019	0.949/47.3
Methods	Cat $\mapsto$ Dog			Dog $\mapsto$ Cat		
	AL $\uparrow$ /Rg $\uparrow$	RL $\uparrow$ /Sm $\downarrow$	ACC $\uparrow$ /FID $\downarrow$	AL $\uparrow$ /Rg $\uparrow$	RL $\uparrow$ /Sm $\downarrow$	ACC $\uparrow$ /FID $\downarrow$
StarGANv2-Rdm [13]	0.872/0.279	0.979/0.299	0.924/33.9	0.840/0.293	0.974/0.295	0.980/25.7
StarGANv2-Cent [13]	0.738/0.242	0.980/0.284	0.906/50.2	0.742/0.263	0.974/0.282	0.991/38.0
Liu <i>et al.</i> -Rdm [44]	0.930/0.281	0.967/0.174	0.956/37.1	0.933/0.313	0.956/0.216	0.997/27.5
SAVI2I-Rdm [51]	0.973/0.294	0.990/0.142	0.942/28.6	0.968/0.292	0.992/0.138	0.982/24.2
SAVI2I-Cent [51]	0.955/0.235	0.988/0.141	0.976/53.6	0.957/0.251	0.995/0.132	0.997/31.4
MonoPix (Ours)	0.986/0.398	0.990/0.050	0.970/41.1	0.983/0.407	0.984/0.076	0.997/14.4

evaluations can be found in the supplementary material. It is also noteworthy to mention that interpolating towards the style centroid leads to a smoother and deterministic translation (variant “-Cent”), yet taking such a trajectory can cause duplicated or unrealistic patterns for style-guided approaches.

#### 4.4 Applications in Low-Level Image Processing

In addition to cross-domain continuous translation, we show how MonoPix, along with the inference strategies can be used to facilitate low-level tasks.

**Low-Light Image Enhancement.** We integrate the proposed contrastive modulation approach with EnlightenGAN [28], a previous unsupervised one-to-one mapping (US-O2O) method in this area, and compare the performance on LOL test set with several state-of-the-art methods, including supervised one-to-one mapping (SL-O2O) method TBEFN [47], and reference-guided supervised (SL-Ref) approach KinD++ [78]. For fairness, we retrain EnlightenGAN on LOL training set, which yields an improved performance. We choose PSNR to guide the ternary search for this task unless specified.

To figure out the potential of our controllable generation scheme, we first implement exhaustive inference (EI), whose results can be viewed as the upper bound. As can be found in Table 2, MonoPix yields an average PSNR of 21.65, which is considerably higher than the one-to-one mapping baselines. This is because a rigid O2O mapping is not sufficient to model the diversity of data distribution. Interestingly, this result is even comparable to state-of-the-art reference-guided supervised approach KinD++, but note that MonoPix can produce competitive results without any expert knowledge or paired supervision, and streamlines the enhancement pipeline notably. The superiority of MonoPix is also supported from SSIM and LPIPS scores. When using a 7-times ternary search, we observe a tiny and acceptable performance drop. This evidence not only validates the monotonicity and continuity of the shaped generation manifold, but also demonstrates the effectiveness of our inference strategy. Next, we formulate a non-reference-guided variant by training an expert network with 10% of paired supervision (during training the backbone conditional generator is frozen). Shown in the experiment named “LP”, the model then degrades to an O2O mapping-based translator where the previous unsupervised contrastive training stage now serves as a pre-text task. Though not competitive as the

Table 2: Quantitative results on LOL test set. ELGAN stands for EnlightenGAN

Method	Type	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
TBEFN [47]	SL-O2O	17.35	0.777	0.210
KinD++ [78]	SL-Ref	21.80	0.829	0.158
KinD++ [78]	SL-O2O	17.75	0.758	0.198
ELGAN [28]	US-O2O	19.01	0.709	0.274
Ours-EI	US-Ref	21.65	0.731	0.273
Ours-TS	US-Ref	21.55	0.731	0.273
Ours-LP	Few Shot	19.75	0.725	0.272

Table 3: Quantitative results on SIDD natural noise generation

Method	Type	AKLD $\downarrow$
CBDNet [20]		0.728
ULRD [4]	SL	0.545
GRDN [32]		0.443
DANet [74]		0.212
Ours-EI		0.137
Ours-TS	US	0.139

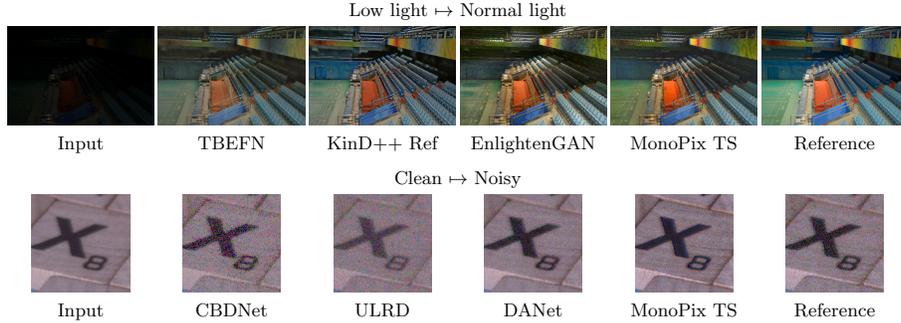


Fig. 5: Visual comparisons on the LOL low-light enhancement task and SIDD natural noise generation



Fig. 6: Examples on fine-grained pixel-level control. From left to right, we provide the input image, spatial control signal, and the modulation result respectively

ternary search variant, it still produces competitive results, yet with a complexity of only  $O(1)$ . Visual comparisons on this task can be found in Figure 5.

**Blind Image Noise Generation.** Different from the commonly discussed task of image denoising, noise generation takes the reverse consideration, and typically is an O2M mapping task. We show in this part the potential and benefits of introducing the proposed continuous modulation approach into this field, and compare with state-of-the-art methods CBDNet [20], ULRD [4], GRDN [32], and DANet [74]. Quantitative and visual results can be found in Table 3 and Figure 5 respectively, where our model (“TS” variant is guided by AKLD) again exhibits its superiority and generalization ability. To the best of our knowledge, MonoPix is the first work that proposes to continuously model natural noise in an unsupervised and spatially-controlled manner. More visual comparisons can be found in the supplementary material.

**Pixel-level Monotonic Control.** In all the previous experiments, we change the overall translation strength by assigning each pixel with the same intensity. Recall that a fine-grained pixel-level control can be naturally obtained by customizing  $c^{i,j}$ , we here provide a few examples in Figure 6. Note that the linearity of manifold further enables a slightly out-of-the-bound inference, especially for uni-directional tasks like illumination and noise control.

## 5 Discussion

### 5.1 Ablation Studies

To further examine the impact of the introduced method, we conduct ablation studies on Yosemite summer-winter translation with two key components that construct our approach: the monotonicity loss (carried along with the contrastive training scheme) and domain fidelity loss. Shown in Table 4 (a) and Figure 7, the variant with none of these losses degrades to an O2O mapping, as we have expected. Note that the AL score is near zero since Pearson correlation coefficient can be negative. When we further add the monotonicity loss, the model succeeds in catching confidence differences and produces a coarse yet continuous translation. Our full model, as shown in variant (c), yields a better fidelity and wider control range. We find the translation quality is also improved under this setting, which is because fidelity loss also strengthens the translation monotonicity and encourages to explore more domain-specific features.

### 5.2 Sensitivity Analysis

We provide the sensitivity analysis on the hyperparameter  $\epsilon$  in  $\mathcal{L}_{mono}$  and  $\mathcal{L}_{df}$ , robustness of guiding criterion in ternary search, and the impact of inference complexity in this part. For better comparisons, we label the variants with varied  $\epsilon$  consecutively following Table 4(c). The results in experiment (d) illustrate that merely asking the discriminator to produce a higher confidence level is insufficient and easily collapses to a trivial solution. In contrast, if we set  $\epsilon$  to 1.0, as in experiment (e), the whole loss function is then largely dominated by  $\mathcal{L}_{mono}$  and  $\mathcal{L}_{df}$ , which ultimately leads to an unpleasant translation. Here we note that variant (e) achieves a better FID score than (d), as the latter suffers from a severe mode collapse issue; however considering the unrealistic translation results, setting a high  $\epsilon$  inevitably introduces a quality degradation for our full model (c). Concluded from these experiments, hyperparameter  $\epsilon$  provides a balance for control intensity and visual quality, and is suggested to increase from a small value. Typically, setting  $\epsilon = 0.5$  can yield acceptable translations for most of the examined tasks. Figure 8 shows that MonoPix can always produce pleasant results and correspondingly achieves the best score when

Table 4: Ablation studies on Yosemite dataset

Variant	$\mathcal{L}_{mono}$	$\mathcal{L}_{df}$	$\epsilon$	Summer $\rightarrow$ Winter			Winter $\rightarrow$ Summer		
				AL $\uparrow$ /Rg $\uparrow$	RL $\uparrow$ /Sm $\downarrow$	ACC $\uparrow$ /FID $\downarrow$	AL $\uparrow$ /Rg $\uparrow$	RL $\uparrow$ /Sm $\downarrow$	ACC $\uparrow$ /FID $\downarrow$
(a)				0.143/0.000	0.999/0.000	0.733/49.9	0.105/0.000	1.000/0.000	0.702/60.0
(b)	$\checkmark$		0.5	0.927/0.091	0.914/0.008	0.866/41.3	0.963/0.097	0.964/0.010	0.847/50.3
(c)	$\checkmark$	$\checkmark$		0.939/0.154	0.940/0.015	0.950/38.1	0.960/0.145	0.964/0.019	0.949/47.3
(d)	$\checkmark$	$\checkmark$	0.0	0.415/0.000	1.000/0.000	0.735/50.0	0.250/0.000	1.000/0.000	0.712/59.5
(e)	$\checkmark$	$\checkmark$	1.0	0.771/0.569	0.861/0.210	0.996/39.7	0.774/0.571	0.873/0.204	0.994/48.7



Fig. 7: Visual comparisons on the ablation study and ternary search criterion sensitivity analysis. We label for each translated im- (top) and impact of age the LPIPS score against the original input to ternary search complexity (bottom) provide a better comparison

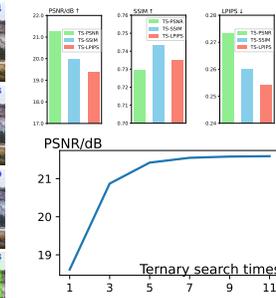


Fig. 8: Sensitivity analysis on the choice of ternary search criterion and impact of age the LPIPS score against the original input to ternary search complexity (bottom) provide a better comparison

guided with a specific criterion. Though not surprising, these results also provide a strong evidence on the feasibility and robustness of our proposals. We have also tested the impact of  $N$  which directly controls the inference complexity, where a 5 to 7 times of ternary search is found sufficient to provide competitive results.

## 6 Conclusions

We investigate in this paper an interesting yet important task - unsupervised continuous pixel-level modulation, and present MonoPix, the first successful solution by devising a novel contrastive modulation framework and the corresponding constraints. MonoPix can be deployed with selective inference complexity, and achieves a state-of-the-art performance on both the translation continuity and quality. The proposed contrastive modulation approach also provides a new alternative to paired or one-to-one mapping pipelines in low-level visions.

We also note that there are still several challenges under the current MonoPix framework, such as pixel-level multi-modal and multi-attribute control, absolute-relative linearity balance, better constraints and evaluation metrics etc. We hope the proposed method could attract more research interest in continuous pixel-level modulation, and help to shape our view on non-O2O mappings.

**Acknowledgment.** This work was supported in part by the National Natural Science Foundation of China under Grants 61731002 and 62071425, in part by the Zhejiang Key Research and Development Plan under Grants 2019C01002 and 2019C03131, in part by Huawei Cooperation Project, and in part by the Zhejiang Provincial Natural Science Foundation of China under Grant LY20F010016.

## References

1. Abdelhamed, A., Lin, S., Brown, M.S.: A high-quality denoising dataset for smart-phone cameras. In: CVPR. pp. 1692–1700 (2018)
2. Ali, J., Lucy, C., Phillip, I.: On the “steerability” of generative adversarial networks. In: ICLR (2020)
3. Almahairi, A., Rajeshwar, S., Sordoni, A., Bachman, P., Courville, A.: Augmented CycleGAN: Learning many-to-many mappings from unpaired data. In: ICML. pp. 195–204 (2018)
4. Brooks, T., Mildenhall, B., Xue, T., Chen, J., Sharlet, D., Barron, J.T.: Unprocessing images for learned raw denoising. In: CVPR. pp. 11036–11045 (2019)
5. Cai, H., He, J., Qiao, Y., Dong, C.: Toward interactive modulation for photo-realistic image restoration. In: CVPR. pp. 294–303 (2021)
6. Chen, C., Chen, Q., Xu, J., Koltun, V.: Learning to see in the dark. In: CVPR. pp. 3291–3300 (2018)
7. Chen, D., Fan, Q., Liao, J., Aviles-Rivero, A., Yuan, L., Yu, N., Hua, G.: Controllable image processing via adaptive filterbank pyramid. *TIP* **29**, 8043–8054 (2020)
8. Chen, D., Yuan, L., Liao, J., Yu, N., Hua, G.: StyleBank: An explicit representation for neural image style transfer. In: CVPR. pp. 1897–1906 (2017)
9. Chen, J., Chen, J., Chao, H., Yang, M.: Image blind denoising with generative adversarial network based noise modeling. In: CVPR. pp. 3155–3164 (2018)
10. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: ICML. pp. 1597–1607 (2020)
11. Chen, Y.C., Xu, X., Tian, Z., Jia, J.: Homomorphic latent space interpolation for unpaired image-to-image translation. In: CVPR. pp. 2408–2416 (2019)
12. Choi, Y., Choi, M., Kim, M., Ha, J.W., Kim, S., Choo, J.: StarGAN: Unified generative adversarial networks for multi-domain image-to-image translation. In: CVPR. pp. 8789–8797 (2018)
13. Choi, Y., Uh, Y., Yoo, J., Ha, J.W.: StarGAN v2: Diverse image synthesis for multiple domains. In: CVPR. pp. 8188–8197 (2020)
14. Collins, E., Bala, R., Price, B., Susstrunk, S.: Editing in style: Uncovering the local semantics of GANs. In: CVPR. pp. 5771–5780 (2020)
15. Dumoulin, V., Shlens, J., Kudlur, M.: A learned representation for artistic style. In: ICLR (2017)
16. Gatys, L., Ecker, A.S., Bethge, M.: Texture synthesis using convolutional neural networks. In: *NeurIPS*. vol. 28 (2015)
17. Girshick, R.: Fast R-CNN. In: ICCV (2015)
18. Gong, R., Li, W., Chen, Y., Gool, L.V.: DLOW: Domain flow for adaptation and generalization. In: CVPR. pp. 2477–2486 (2019)
19. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: *NeurIPS*. vol. 27 (2014)
20. Guo, S., Yan, Z., Zhang, K., Zuo, W., Zhang, L.: Toward convolutional blind denoising of real photographs. In: CVPR. pp. 1712–1722 (2019)
21. He, J., Dong, C., Qiao, Y.: Modulating image restoration with continual levels via adaptive feature modification layers. In: CVPR. pp. 11056–11064 (2019)
22. He, J., Dong, C., Qiao, Y.: Interactive multi-dimension modulation with dynamic controllable residual learning for image restoration. In: ECCV. pp. 53–68 (2020)
23. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)

24. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: GANs trained by a two time-scale update rule converge to a local nash equilibrium. In: *NeurIPS*. vol. 30 (2017)
25. Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: *ICCV*. pp. 1501–1510 (2017)
26. Huang, X., Liu, M.Y., Belongie, S., Kautz, J.: Multimodal unsupervised image-to-image translation. In: *ECCV*. pp. 172–189 (2018)
27. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: *CVPR*. pp. 1125–1134 (2017)
28. Jiang, Y., Gong, X., Liu, D., Cheng, Y., Fang, C., Shen, X., Yang, J., Zhou, P., Wang, Z.: EnlightenGAN: Deep light enhancement without paired supervision. *TIP* **30**, 2340–2349 (2021)
29. Jo, S.Y., Lee, S., Ahn, N., Kang, S.J.: Deep arbitrary HDRI: Inverse tone mapping with controllable exposure changes. *TMM* (2021)
30. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: *ECCV*. pp. 694–711 (2016)
31. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: *CVPR*. pp. 4401–4410 (2019)
32. Kim, D.W., Ryun Chung, J., Jung, S.W.: GRDN: Grouped residual dense network for real image denoising and GAN-based real-world noise modeling. In: *CVPRW* (2019)
33. Kim, H., Choi, Y., Kim, J., Yoo, S., Uh, Y.: Exploiting spatial dimensions of latent in GAN for real-time image editing. In: *CVPR*. pp. 852–861 (2021)
34. Klyuchka, M., Neiterman, E.H., Ben-Artzi, G.: CEL-Net: Continuous exposure for extreme low-light imaging. *arXiv:2012.04112* (2020)
35. Lee, C.H., Liu, Z., Wu, L., Luo, P.: MaskGAN: Towards diverse and interactive facial image manipulation. In: *CVPR*. pp. 5549–5558 (2020)
36. Lee, H.Y., Tseng, H.Y., Huang, J.B., Singh, M., Yang, M.H.: Diverse image-to-image translation via disentangled representations. In: *ECCV*. pp. 35–51 (2018)
37. Lee, H., Kim, T., Son, H., Baek, S., Cheon, M., Lee, S.: Smoother network tuning and interpolation for continuous-level image processing. *arXiv:2010.02270* (2020)
38. Li, C., Guo, C., Han, L.H., Jiang, J., Cheng, M.M., Gu, J., Loy, C.C.: Low-light image and video enhancement using deep learning: a survey. *TPAMI* (2021)
39. Li, C., Wand, M.: Combining Markov random fields and convolutional neural networks for image synthesis. In: *CVPR*. pp. 2479–2486 (2016)
40. Li, Y., Fang, C., Yang, J., Wang, Z., Lu, X., Yang, M.H.: Diversified texture synthesis with feed-forward networks. In: *CVPR*. pp. 3920–3928 (2017)
41. Li, Y., Fang, C., Yang, J., Wang, Z., Lu, X., Yang, M.H.: Universal style transfer via feature transforms. In: *NeurIPS*. vol. 30 (2017)
42. Ling, H., Kreis, K., Li, D., Kim, S.W., Torralba, A., Fidler, S.: EditGAN: High-precision semantic image editing. In: *NeurIPS*. vol. 34 (2021)
43. Lira, W., Merz, J., Ritchie, D., Cohen-Or, D., Zhang, H.: GANHopper: Multi-hop GAN for unsupervised image-to-image translation. In: *ECCV*. pp. 363–379 (2020)
44. Liu, Y., Sangineto, E., Chen, Y., Bao, L., Zhang, H., Sebe, N., Lepri, B., Wang, W., De Nadai, M.: Smoothing the disentangled latent style space for unsupervised image-to-image translation. In: *CVPR*. pp. 10785–10794 (2021)
45. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: *CVPR*. pp. 3431–3440 (2015)
46. Lore, K.G., Akintayo, A., Sarkar, S.: LLNet: A deep autoencoder approach to natural low-light image enhancement. *PR* **61**, 650–662 (2017)

47. Lu, K., Zhang, L.: TBEFN: A two-branch exposure-fusion network for low-light image enhancement. *TMM* **23**, 4093–4105 (2020)
48. Lu, Y., Tai, Y.W., Tang, C.K.: Attribute-guided face generation using conditional CycleGAN. In: *ECCV*. pp. 282–297 (2018)
49. Lv, Z., Li, X., Li, X., Li, F., Lin, T., He, D., Zuo, W.: Learning semantic person image generation by region-adaptive normalization. In: *CVPR*. pp. 10806–10815 (2021)
50. Maas, A.L., Hannun, A.Y., Ng, A.Y., et al.: Rectifier nonlinearities improve neural network acoustic models. In: *ICML*. vol. 30, p. 3 (2013)
51. Mao, Q., Tseng, H.Y., Lee, H.Y., Huang, J.B., Ma, S., Yang, M.H.: Continuous and diverse image-to-image translation via signed attribute vectors. *IJCV* **130**, 517–549 (2022)
52. Mao, X., Li, Q., Xie, H., Lau, R.Y., Wang, Z., Paul Smolley, S.: Least squares generative adversarial networks. In: *ICCV*. pp. 2794–2802 (2017)
53. Ni, S., Cao, X., Yue, T., Hu, X.: Controlling the rain: From removal to rendering. In: *CVPR*. pp. 6328–6337 (2021)
54. Park, T., Efros, A.A., Zhang, R., Zhu, J.Y.: Contrastive learning for unpaired image-to-image translation. In: *ECCV*. pp. 319–345 (2020)
55. Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: *CVPR*. pp. 2337–2346 (2019)
56. Pizzati, F., Cerri, P., de Charette, R.: CoMoGAN: continuous model-guided image-to-image translation. In: *CVPR*. pp. 14288–14298 (2021)
57. Pumarola, A., Agudo, A., Martinez, A.M., Sanfeliu, A., Moreno-Noguer, F.: GAN-imation: Anatomically-aware facial animation from a single image. In: *ECCV*. pp. 818–833 (2018)
58. Ren, D., Zuo, W., Hu, Q., Zhu, P., Meng, D.: Progressive image deraining networks: A better and simpler baseline. In: *CVPR*. pp. 3937–3946 (2019)
59. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *MICCAI*. pp. 234–241 (2015)
60. Shen, Y., Zhou, B.: Closed-form factorization of latent semantics in GANs. In: *CVPR*. pp. 1532–1540 (2021)
61. Shi, J., Xu, N., Xu, Y., Bui, T., Dernoncourt, F., Xu, C.: Learning by planning: Language-guided global image editing. In: *CVPR*. pp. 13590–13599 (2021)
62. Shoshan, A., Mechrez, R., Zelnik-Manor, L.: Dynamic-Net: Tuning the objective without re-training for synthesis tasks. In: *ICCV*. pp. 3215–3223 (2019)
63. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556* (2014)
64. Sun, X., Li, M., He, T., Fan, L.: Enhance images as you like with unpaired learning. *arXiv:2110.01161* (2021)
65. Ulyanov, D., Vedaldi, A., Lempitsky, V.: Instance normalization: The missing ingredient for fast stylization. *arXiv:1607.08022* (2016)
66. Upchurch, P., Gardner, J., Pleiss, G., Pless, R., Snively, N., Bala, K., Weinberger, K.: Deep feature interpolation for image content changes. In: *CVPR*. pp. 7064–7073 (2017)
67. Wang, W., Guo, R., Tian, Y., Yang, W.: CFSNet: Toward a controllable feature space for image restoration. In: *ICCV*. pp. 4140–4149 (2019)
68. Wang, X., Yu, K., Dong, C., Tang, X., Loy, C.C.: Deep network interpolation for continuous imagery effect transition. In: *CVPR*. pp. 1692–1701 (2019)
69. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *TIP* **13**(4), 600–612 (2004)

70. Wei, C., Wang, W., Yang, W., Liu, J.: Deep Retinex decomposition for low-light enhancement. In: BMVC (2018)
71. Wu, P.W., Lin, Y.J., Chang, C.H., Chang, E.Y., Liao, S.W.: RelGAN: Multi-domain image-to-image translation via relative attributes. In: ICCV. pp. 5914–5922 (2019)
72. Ye, M., Zhang, X., Yuen, P.C., Chang, S.F.: Unsupervised embedding learning via invariant and spreading instance feature. In: CVPR. pp. 6210–6219 (2019)
73. Yu, X., Chen, Y., Liu, S., Li, T., Li, G.: Multi-mapping image-to-image translation via learning disentanglement. In: NeurIPS. vol. 32 (2019)
74. Yue, Z., Zhao, Q., Zhang, L., Meng, D.: Dual adversarial network: Toward real-world noise removal and noise generation. In: ECCV. pp. 41–58 (2020)
75. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H., Shao, L.: Multi-stage progressive image restoration. In: CVPR. pp. 14821–14831 (2021)
76. Zhang, H., Dana, K.: Multi-style generative network for real-time transfer. In: ECCVW (2018)
77. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR. pp. 586–595 (2018)
78. Zhang, Y., Guo, X., Ma, J., Liu, W., Zhang, J.: Beyond brightening low-light images. *IJCV* **129**(4), 1013–1037 (2021)
79. Zhu, J., Shen, Y., Zhao, D., Zhou, B.: In-domain GAN inversion for real image editing. In: ECCV. pp. 592–608 (2020)
80. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: ICCV. pp. 2223–2232 (2017)
81. Zhu, P., Abdal, R., Qin, Y., Wonka, P.: SEAN: Image synthesis with semantic region-adaptive normalization. In: CVPR. pp. 5104–5113 (2020)
82. Zhuang, P., Koyejo, O., Schwing, A.G.: Enjoy your editing: Controllable GANs for image editing via latent space navigation. In: ICLR (2021)