Learn from Unpaired Data for Image Restoration: A Variational Bayes Approach

Dihan Zheng, Xiaowen Zhang, Kaisheng Ma, Chenglong Bao*

Abstract—Collecting paired training data is difficult in practice, but the unpaired samples broadly exist. Current approaches aim at generating synthesized training data from the unpaired samples by exploring the relationship between the corrupted and clean data. This work proposes LUD-VAE, a deep generative method to learn the joint probability density function from data sampled from marginal distributions. Our approach is based on a carefully designed probabilistic graphical model in which the clean and corrupted data domains are conditionally independent. Using variational inference, we maximize the evidence lower bound (ELBO) to estimate the joint probability density function. Furthermore, we show that the ELBO is computable without paired samples under the inference invariant assumption. This property provides the mathematical rationale of our approach in the unpaired setting. Finally, we apply our method to real-world image denoising and super-resolution tasks and train the models using the synthetic data generated by the LUD-VAE. Experimental results validate the advantages of our method over other learnable approaches.

Index Terms—Image restoration, unpaired degradation modeling, graphical model, variational auto-encoder.

1 INTRODUCTION

Image restoration aims to recover the underlying clean image \mathbf{x} from the corrupted observation \mathbf{y} ,

$$\mathbf{y} = \mathcal{T}(\mathbf{x}) + \mathbf{n},\tag{1}$$

where **n** represents the noise, and \mathcal{T} represents the degradation operation. This task is one of the fundamental problems in computer vision and has been extensively studied for decades [3], [4], [5]. In recent years, deep learning has achieved astonishing success in image restoration problems, such as image denoising [6], [7], [8] and superresolution [9], [10], [11], [12]. However, the success of these methods requires large quantities of paired training data, and the restoration performance is sensitive to the degradation types [13], [14], [15]. For example, one Gaussian denoising network usually performs poorly for real-world noisy images due to the noise discrepancy between Gaussian noise and real-world noise [16]. Meanwhile, collecting paired training data for real-world image restoration is cumbersome and expensive due to the complex camera image signal processing (ISP) pipeline [17], [18], [19]. The aforementioned problems make real-world image restoration a challenging task. On the other hand, unpaired data broadly exists and is easily accessible in many situations. For example, it is easy to obtain many images of different resolutions or noisy and clean images through the internet [20]. Consequently, designing deep learning methods with unpaired data is of significant research importance and deserves deep exploration.

Currently, two common strategies are along this line. One is the unsupervised image restoration methods with a single corrupted image [16], [21], [22] or corrupted image dataset [23], [24], which do not take clean images into consideration. As a result, those methods are either timeconsuming or inferior to supervised methods. The other strategy is to learn the degradation model from the unpaired datasets. After learning a generative model to construct the synthetic paired training data, it trains an image restoration model using conventional supervised deep learning methods. The main difficulty of these methods is to develop effective methods so that the synthetic paired data is close to the underlying paired data. Current methods mainly adopt generative adversarial networks (GANs) [25] and draw on the idea from cycle-consistency constraint in Cycle-GAN [26], [27], [28] and domain adversarial training [29], [30], [31]. However, these methods often require careful adjustment of different losses, and the heuristic constraint of cycle consistency is too weak for this problem and lacks theoretical rigorousness [1], [32]. More importantly, those GAN-based methods usually obtain a deterministic mapping while ignoring the randomness in the degradation generation process. Recently, the DeFlow [1] method, which models the unpaired degradation process using a conditional flow model, has shown promising performance in super-resolution. In Figure 1, it shows that the noise distribution generated by Deflow does not well match the true noise distribution.

In this work, we propose LUD-VAE, a variational autoencoder (VAE) based degradation modeling method for unpaired training data. Mathematically, given $\mathbf{x} \sim p(\mathbf{x})$ and $\mathbf{y} \sim p(\mathbf{y})$, our goal is to approximate the joint distribution $p(\mathbf{x}, \mathbf{y})$. In general, as shown in [1], this task is difficult and

Dihan Zheng is with the Yau Mathematical Sciences Center, Tsinghua University, Beijing 100084, China (e-mail: zhengdh19@mails.tsinghua.edu.cn)

Xiaowen Zhang is with the Hisilicon, Shanghai 300060, China (e-mail: zhangxiaowen9@hisilicon.com)

Kaisheng Ma is with the Institute for Interdisciplinary Information Science, Tsinghua University, Beijing 100084, China (e-mail: kaisheng@mail.tsinghua.edu.cn)

Chenglong Bao is with the Yau Mathematical Sciences Center, Tsinghua University, Beijing 100084, China, and Yanqi Lake Beijing Institute of Mathematical Sciences and Applications, Beijing 101408, China (e-mail: clbao@mail.tsinghua.edu.cn)

^{• *} Corresponding author.

Image: Clean imageReal noiseNoise from DeFlow [1]Noise from LUD-VAEImage: Clean image

Fig. 1: Synthetic noisy images obtained by DeFlow and LUD-VAE learned from unpaired noisy-clean dataset AIM2019 [2]. Patches of 32×32 from each images are chosen for comparison, the noise distribution in these patches are shown.

has no unique solution. Thus, it motivates us to carefully design a computable generative graph that well approximates a reasonable solution. More specifically, this graph consists of two independent latent variables: \mathbf{z} , which encodes the image information, and \mathbf{z}_{n} , which encodes the degradation information. The generation relationship is designed as follows: x is generated from z; y is generated from z and z_n . Using the idea from VAE, we introduce an encoder network for the inference and a decoder network for the generative process so that we can model the conditional distribution $p(\mathbf{y}|\mathbf{x})$. In addition, we impose the inference invariant condition on \mathbf{z} that requires the same latent representations of paired x and y and show that the ELBO can be computed via unpaired corrupted and clean data. This property gives us a transparent and explainable loss for training the networks. The Figure 1 demonstrates that the proposed method can learn the noise distribution accurately. Finally, we apply the LUD-VAE model to the problem of real-world image denoising and super-resolution. LUD-VAE is used to learn the degradation model with unpaired data and synthesize paired training data for the downstream supervised models. We test the performance of LUD-VAE on two real-world super-resolution datasets: AIM2019 [2] and NTIRE2020 [33], and one real-world image denoising dataset: SIDD [13]. Experimental results show that the proposed LUD-VAE model outperforms GAN-based approaches. Compared with the recent flow-based method, our model achieves comparable results with much fewer parameters. Our main contributions are summarized as follows.

• We propose a degradation modeling method, called the LUD-VAE, by constructing a new generative graph. Using the VAE framework, we derive an explainable loss function that decouples the dependency between noisy and clean images. The work is the first attempt to model the underlying paired data distribution for image restoration from the VAE perspective to the best of our knowledge.

- We design a hierarchical structure for LUD-VAE to effectively learn the degradation process with an economic network. Moreover, the model is selfsupervised and does not require pre-trained networks.
- Experimental results in the task of real-world image denoising and super-resolution validate the advantages of the proposed LUD-VAE.

2 RELATED WORK

Deep image restoration. In recent years, deep learning based image restoration methods have significantly improved performance over traditional methods. In additive white Gaussian noise (AWGN) removal task, DnCNN [6] uses deep convolutional networks to predict noise from the noisy images. Further, FFDNet [7] takes the noise level as network input to deal with images with different noise levels. Recently, researchers have designed different structures to improve denoising performance [34], [35], [36]. While these methods have achieved state-of-the-art performance on the synthetic datasets, they still require numerous paired training data to learn the restoration process. In real-world image restoration tasks, collecting paired data is difficult due to the sophisticated imaging process. Some real-world image restoration datasets such as image denoising [8], [13], [37], [38], [39] and super-resolution [27], [40], [41]. They use statistical methods to synthesize clean images from multiple noisy observations, requiring careful setups and procedures. With these datasets, many methods are proposed for realworld image restoration, such as using various attention modules [42], [43], [44], vision transformers [45], [46], [47], and GANs [11], [48]. There are also some unsupervised or self-supervised methods, which reduce the requirements for training datasets. N2N [49] proposes to use paired noise images to learn an image denoising network. Subsequently,

N2V [50] and N2S [51] further reduce this requirement. They use a "blind spot" network to learn a denoising model directly from the noise data set without clean images. However, these methods need to assume that the image noise is spatially independent, limiting their application to real data. Unpaired degradation modeling. Learning the degradation model from unpaired data can be considered the image-toimage transfer task, a long-standing problem in computer vision. Most of the existing works employ the GANs [25], mainly using cycle-consistency [26], [27] proposed in Cycle-GAN [28] and domain adversarial [29], [30], [31] to characterize the conditional relationship between $p(\mathbf{x})$ and $p(\mathbf{y})$. However, these methods often use heuristic losses, lack theoretical guarantees, and need elaborate fine-tuning of those losses [1]. Meanwhile, GAN-based methods may have some issues, such as unstable training [52] and mode collapse [53]. Our model is based on variational inference, each term in our loss function can be derived from the evidence lower bound, and there are no adversarial losses as in GANs. There are also handcrafted methods to synthesize degradation images [54], but they lack generalization ability [1], [14]. Recently, [1] proposed the DeFlow model, a flow-based degradation modeling method without paired data, which has achieved excellent performance on real-world superresolution tasks.

Unpaired learning with VAEs. Variational auto-encoder (VAE) [55] based unpaired learning for degradation modeling is currently less developed; here, we investigate some related topics. [56] proposed to train an Energy-Based Model (EBM) in the latent space of a trained VAE to realize imageto-image transfer. However, this method uses Markov Chain Monte Carlo (MCMC) algorithms to sample from the latent space, which leads to slow generation speed. Meanwhile, this method has no theoretical constraints to ensure the rationality of the transformation of \mathbf{x} to \mathbf{y} , which is mainly used for the transformation between human faces or animal images. In [16], a single image based unsupervised denoising method with VAEs is proposed, but this method requires training a new network for each image. [57] proposed a dataset based unsupervised denoising algorithm for microscopy images with VAEs, which needs to know the noise distribution and requires multiple re-samplings to boost the performance. In addition, it only verifies the effectiveness for microscopy images and text images, whereas the texture and content of natural images are more complicated. In this work, we learn the degradation model with unpaired data and design an explainable and reasonable loss with successful applications in real-world image restoration tasks.

3 OUR METHODOLOGY

In this section, we present our method for learning the unknown degradation model using unpaired noisy and clean images. Formally, assume the data $\{\mathbf{x}_i\}$ and $\{\mathbf{y}_j\}$ are i.i.d. sampled from $p(\mathbf{x})$ and $p(\mathbf{y})$ respectively, and our goal is to generate paired samples from the conditional distribution $p(\mathbf{y}|\mathbf{x})$. In the following context, we assume clean images lie in the source domain and noisy images lie in the target domain.



Fig. 2: Graphical model of the image generation process. (\mathbf{x}, \mathbf{y}) are paired clean and noisy images; $(\mathbf{z}, \mathbf{z}_n)$ are latent variables for generating (\mathbf{x}, \mathbf{y}) .

3.1 Basic idea

To find the transformation from x to y, one straightforward idea is to estimate the conditional density $p(\mathbf{y}|\mathbf{x})$. However, it is difficult to model the conditional density function directly due to the lack of paired data. Instead, we consider to model the joint density function $p(\mathbf{x}, \mathbf{y})$ in this work. To leverage the information of the unpaired data, our basic idea is to decouple the joint density function into the source domain and target domain. Since the independent assumption for x, y does not hold in practice, we impose the conditional independence by assuming the joint random variable (\mathbf{x}, \mathbf{y}) has two latent variables: \mathbf{z} and $\mathbf{z}_{\mathbf{n}}$. For a paired data (\mathbf{x}, \mathbf{y}) sampled from $p(\mathbf{x}, \mathbf{y})$, we assume that the image content and the degradation information are from latent variables z and z_n respectively. See Figure 2 for the generative graph. Under the above assumptions, the conditional joint distribution becomes

$$p(\mathbf{x}, \mathbf{y} | \mathbf{z}, \mathbf{z}_{\mathbf{n}}) = p(\mathbf{x} | \mathbf{z}) p(\mathbf{y} | \mathbf{z}, \mathbf{z}_{\mathbf{n}}),$$
(2)

and the conditional log-likelihood is

$$\sum_{i} \log p(\mathbf{x}_{i}, \mathbf{y}_{i} | \mathbf{z}, \mathbf{z}_{n})$$

$$= \sum_{i} \log p(\mathbf{x}_{i} | \mathbf{z}) + \log p(\mathbf{y}_{i} | \mathbf{z}, \mathbf{z}_{n})$$

$$= \sum_{i} \log p(\mathbf{x}_{i} | \mathbf{z}) + \sum_{j} \log p(\mathbf{y}_{j} | \mathbf{z}, \mathbf{z}_{n}),$$
(3)

which removes the dependence of paired data. Inspired by the above observation, we can approximate the joint distribution $p(\mathbf{x}, \mathbf{y})$ using the VAE framework in the next subsection.

3.2 Proposed LUD-VAE method

To estimate the joint density function $p(\mathbf{x}, \mathbf{y})$ with the graphical model given by Figure 2, we apply the variational inference framework. Note that the log-likelihood function $\log p(\mathbf{x}, \mathbf{y})$ has the following decomposition:

$$\log p(\mathbf{x}, \mathbf{y}) = \mathbb{E}_{q(\mathbf{z}, \mathbf{z}_{\mathbf{n}} | \mathbf{x}, \mathbf{y})} \log \frac{p(\mathbf{z}, \mathbf{z}_{\mathbf{n}}, \mathbf{x}, \mathbf{y})}{q(\mathbf{z}, \mathbf{z}_{\mathbf{n}} | \mathbf{x}, \mathbf{y})} + D_{\mathrm{KL}} \left(q(\mathbf{z}, \mathbf{z}_{\mathbf{n}} | \mathbf{x}, \mathbf{y}) \| p(\mathbf{z}, \mathbf{z}_{\mathbf{n}} | \mathbf{x}, \mathbf{y}) \right),$$
(4)

where $D_{\rm KL}$ is the Kullback–Leibler (KL) divergence. Since the second term in (4) is non-negative, the first term in (4) provides a lower bound of the log-likelihood. In fact, the expectation term in (4) is equal to

$$\mathbb{E}_{q(\mathbf{z},\mathbf{z}_{n}|\mathbf{x},\mathbf{y})}\log p(\mathbf{x},\mathbf{y}|\mathbf{z},\mathbf{z}_{n}) - D_{\mathrm{KL}}(q(\mathbf{z},\mathbf{z}_{n}|\mathbf{x},\mathbf{y})||p(\mathbf{z}_{n},\mathbf{z})),$$
(5)



Fig. 3: The hierarchical architecture used in LUD-VAE.

which is called as the evidence lower bound (ELBO). Thus, instead of maximizing the intractable log-likelihood, we maximize the ELBO. Suppose the image information is contained in the paired data (\mathbf{x}, \mathbf{y}) and the degradation information is only contained in the noisy data \mathbf{y} , we choose the inference model which has the decomposition

$$q(\mathbf{z}_{\mathbf{n}}, \mathbf{z} | \mathbf{x}, \mathbf{y}) = q(\mathbf{z} | \mathbf{x}, \mathbf{y}) q(\mathbf{z}_{\mathbf{n}} | \mathbf{y}).$$
(6)

Moreover, the graphical model in Figure 2 gives

$$p(\mathbf{z}_{\mathbf{n}}, \mathbf{z}) = p(\mathbf{z}_{\mathbf{n}})p(\mathbf{z}), \quad p(\mathbf{x}, \mathbf{y}|\mathbf{z}, \mathbf{z}_{\mathbf{n}}) = p(\mathbf{x}|\mathbf{z})p(\mathbf{y}|\mathbf{z}, \mathbf{z}_{\mathbf{n}}).$$
(7)

Combining (6) with (7), the ELBO given in (5) satisfies

$$ELBO = \underbrace{\mathbb{E}_{q(\mathbf{z}|\mathbf{x},\mathbf{y})} \log p(\mathbf{x}|\mathbf{z}) + \mathbb{E}_{q(\mathbf{z}|\mathbf{x},\mathbf{y})q(\mathbf{z}_{n}|\mathbf{y})} \log p(\mathbf{y}|\mathbf{z}_{n},\mathbf{z})}_{\text{Reconstruction}} - \underbrace{D_{\text{KL}}(q(\mathbf{z}|\mathbf{x},\mathbf{y})||p(\mathbf{z})) - D_{\text{KL}}(q(\mathbf{z}_{n}|\mathbf{y})||p(\mathbf{z}_{n}))}_{\text{KL}}.$$

Due to the existence of $q(\mathbf{z}|\mathbf{x}, \mathbf{y})$ in ELBO, it still needs the paired information. To further decouple this relationship, we define the *inference invariant condition* as

$$q(\mathbf{z}|\mathbf{x}) = q(\mathbf{z}|\mathbf{y}), \quad \forall (\mathbf{x}, \mathbf{y}) \sim p(\mathbf{x}, \mathbf{y}).$$
 (8)

The above condition means that for paired data (x, y), the latent image information z can be obtained from either clean image x or noisy image y. In other words, the latent code z represents the common features between the source domain and target domain. In practice, this condition can be satisfied using a pre-trained network or predefined operations. More discussions related to this condition is present in Section 3.3.

Under the inference invariant condition (8) and setting the inference model $q(\mathbf{z}|\mathbf{x},\mathbf{y})$ to

$$q(\mathbf{z}|\mathbf{x}, \mathbf{y}) := q(\mathbf{z}|\mathbf{x}) = q(\mathbf{z}|\mathbf{y}), \tag{9}$$

4

we have $ELBO = ELBO_{\mathbf{x}} + ELBO_{\mathbf{y}}$ where

$$\text{ELBO}_{\mathbf{x}} := \mathbb{E}_{q(\mathbf{z}|\mathbf{x})} \log p(\mathbf{x}|\mathbf{z}) - \frac{1}{2} D_{\text{KL}}(q(\mathbf{z}|\mathbf{x}) \| p(\mathbf{z})), \quad (10)$$

$$ELBO_{\mathbf{y}} := \mathbb{E}_{q(\mathbf{z}|\mathbf{y})q(\mathbf{z}_{\mathbf{n}}|\mathbf{y})} \log p(\mathbf{y}|\mathbf{z}_{\mathbf{n}}, \mathbf{z})$$
(11)
$$-\frac{1}{2}D_{\mathrm{KL}}(q(\mathbf{z}|\mathbf{y})||p(\mathbf{z})) - D_{\mathrm{KL}}(q(\mathbf{z}_{\mathbf{n}}|\mathbf{y})||p(\mathbf{z}_{\mathbf{n}})).$$

It is clear that $ELBO_x$ is only related to clean images x and $ELBO_y$ is only related to noisy images y. Thus, no paired data is needed in the above formulation. In summary, the loss of the proposed LUD-VAE model is

$$\text{Fotal Loss} = -(\text{ELBO}_{\mathbf{x}} + \text{ELBO}_{\mathbf{y}}).$$

Remark 1. The recent DeFlow model [1] also estimate the joint density $p(\mathbf{x}, \mathbf{y})$ by directly maximizing two marginal log-likelihood functions $\log p(\mathbf{x})$ and $\log p(\mathbf{y})$. In [1], it assumes that there are two latent variables $\mathbf{z}_{\mathbf{x}}$ and $\mathbf{z}_{\mathbf{y}}$ where $\mathbf{z}_{\mathbf{x}}$ and $\mathbf{z}_{\mathbf{y}}$ are not independent, which is the main difference with our method. If $\mathbf{z}_{\mathbf{y}} = \mathbf{z}_{\mathbf{x}} + \mathbf{u}$ where $\mathbf{u} \sim \mathcal{N}(\mu_u, \Sigma_u)$, the log-likelihood of the joint density contains marginal log-likelihood functions and one extra cross term that requires the clean-noisy image pairs. See the Appendix for the detailed derivation. It is worth mentioning that the DeFlow method introduces a domain invariant function and the conditional flow model. This modified model may help decorrelate the latent representations, but investigating the exact mechanism requires more discussion.

3.3 Method settings

In this subsection, we introduce the details of $ELBO_x$ and $ELBO_y$ by specifying the choices of latent variables z, z_n and the network architecture.

Hierarchical structure. In ELBO_x and ELBO_y, the choices of inference models $q(\mathbf{z}|\mathbf{x}), q(\mathbf{z_n}|\mathbf{y})$ and the prior distributions $p(\mathbf{z}), p(\mathbf{z_n})$ are important for the performance of our model. In classical VAE model [55], it sets $p(\mathbf{z})$ and $p(\mathbf{z_n})$ as the normal distribution that facilitates the computation of KL terms in the loss function. However, this may restrict the expressive ability of the prior distribution, and it is pointed out in [58] that classical VAE models usually suffer from the problem of generating unrealistic and blurry images. To address this problem, we adopt the hierarchical representations of both inference models and prior distribution [59]. Let *N* be the number of layers, we assume

$$\mathbf{z} = (\mathbf{z}^1, \dots, \mathbf{z}^N), \quad \mathbf{z}_n = (\mathbf{z}_n^1, \dots, \mathbf{z}_n^N)$$

Thus, the prior distributions $p(\mathbf{z})$ and $p(\mathbf{z_n})$ are

$$p(\mathbf{z}) = p(\mathbf{z}^{1})p(\mathbf{z}^{2}|\mathbf{z}^{1})\cdots p(\mathbf{z}^{N}|\mathbf{z}^{

$$p(\mathbf{z}_{\mathbf{n}}) = p(\mathbf{z}_{\mathbf{n}}^{1})p(\mathbf{z}_{\mathbf{n}}^{2}|\mathbf{z}_{\mathbf{n}}^{1})\cdots p(\mathbf{z}_{\mathbf{n}}^{N}|\mathbf{z}_{\mathbf{n}}^{
(12)$$$$

where $\mathbf{z}^{<K} = (\mathbf{z}^1, \dots, \mathbf{z}^{K-1})$ and $\mathbf{z}_{\mathbf{n}}^{<K} = (\mathbf{z}_{\mathbf{n}}^1, \dots, \mathbf{z}_{\mathbf{n}}^{K-1})$. Similarly, we impose the same order on the inference models to generate latent variables:

$$q(\mathbf{z}|\mathbf{x}) = q(\mathbf{z}^{1}|\mathbf{x})q(\mathbf{z}^{2}|\mathbf{z}^{1},\mathbf{x})\cdots q(\mathbf{z}^{N}|\mathbf{z}^{< N},\mathbf{x}),$$

$$q(\mathbf{z}_{\mathbf{n}}|\mathbf{y}) = q(\mathbf{z}_{\mathbf{n}}^{1}|\mathbf{y})q(\mathbf{z}_{\mathbf{n}}^{2}|\mathbf{z}_{\mathbf{n}}^{1},\mathbf{y})\cdots q(\mathbf{z}_{\mathbf{n}}^{N}|\mathbf{z}_{\mathbf{n}}^{< N},\mathbf{y}).$$
(13)

Combining (12) and (13), the KL divergence in $ELBO_x$ is

$$D_{\mathrm{KL}}(q(\mathbf{z}|\mathbf{x})||p(\mathbf{z})) = D_{\mathrm{KL}}(q(\mathbf{z}^{1}|\mathbf{x})||p(\mathbf{z}^{1})) + \sum_{l=2}^{N} \mathbb{E}_{q(\mathbf{z}^{< l}|\mathbf{x})} \left[D_{\mathrm{KL}}(q(\mathbf{z}^{l}|\mathbf{z}^{< l},\mathbf{x})||p(\mathbf{z}^{l}|\mathbf{z}^{< l})) \right],$$
(14)

and the two KL terms in ELBOy are

$$D_{\mathrm{KL}}(q(\mathbf{z}|\mathbf{y})||p(\mathbf{z})) = D_{\mathrm{KL}}(q(\mathbf{z}^{1}|\mathbf{y})||p(\mathbf{z}^{1})) + \sum_{l=2}^{N} \mathbb{E}_{q(\mathbf{z}^{< l}|\mathbf{y})} \left[D_{\mathrm{KL}}(q(\mathbf{z}^{l}|\mathbf{z}^{< l},\mathbf{y})||p(\mathbf{z}^{l}|\mathbf{z}^{< l})) \right], \quad (15)$$

$$D_{\mathrm{KL}}(q(\mathbf{z_n}|\mathbf{y})||p(\mathbf{z_n})) = D_{\mathrm{KL}}(q(\mathbf{z_n^l}|\mathbf{y})||p(\mathbf{z_n^l})) + \sum_{l=2}^{N} \mathbb{E}_{q(\mathbf{z_n^{< l}}|\mathbf{y})} \left[D_{\mathrm{KL}}(q(\mathbf{z_n^l}|\mathbf{z_n^{< l}},\mathbf{y})||p(\mathbf{z_n^l}|\mathbf{z_n^{< l}})) \right].$$
(16)

Remark 2. The hierarchical structures on latent variables z and z_n impose the conditional dependence of for different layers. Compared to the classical VAE method, it significantly improves the expressive ability of the prior distributions for modeling the complex degradation in image processing.

The choice of latent space. Besides the hierarchical structure on latent variables \mathbf{z}_n , \mathbf{z} , we specify their distributions for facilitating the computation. For $p(\mathbf{z}_n)$, we assume

$$\begin{split} p(\mathbf{z}_n^l) &= \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad p(\mathbf{z}_n^l | \mathbf{z}_n^{< l}) = \mathcal{N}(\mu_p^l(\mathbf{z}_n^{< l}), \sigma_p^l(\mathbf{z}_n^{< l})), \\ q(\mathbf{z}_n^l | \mathbf{z}_n^{< l}, \mathbf{y}) &= \mathcal{N}(\mu_q^l(\mathbf{z}_n^{< l}, \mathbf{y}), \sigma_q^l(\mathbf{z}_n^{< l}, \mathbf{y})), \end{split}$$

where μ_p^l , σ_p^l , μ_q^l , σ_q^l are encoding neural networks for the degrading part. Thus, the KL divergence in (16) is computable.

For $p(\mathbf{z})$, we assume

$$\begin{split} p(\mathbf{z}^{1}) &= \mathcal{U}(\Omega^{1}), \quad p(\mathbf{z}^{l}|\mathbf{z}^{< l}) = \mathcal{U}(\Omega^{l}), \\ q(\mathbf{z}^{l}|\mathbf{z}^{< l}, \mathbf{x}) &= \delta(e^{l}(\mathbf{z}^{< l}, \mathbf{x})), \quad q(\mathbf{z}^{l}|\mathbf{z}^{< l}, \mathbf{y}) = \delta(e^{l}(\mathbf{z}^{< l}, \mathbf{y})), \end{split}$$

where e^l is the encoding neural network for the image contents, Ω^l is the range of network e^l , $\mathcal{U}(\Omega^l)$ denotes the uniform distribution on Ω^l , $\delta(\cdot)$ denotes the delta distribution. In this case, we have

$$D_{\mathrm{KL}}(q(\mathbf{z}|\mathbf{x})||p(\mathbf{z})) = c, \quad D_{\mathrm{KL}}(q(\mathbf{z}|\mathbf{y})||p(\mathbf{z})) = c, \quad (17)$$

where c is a constant. Moreover, let $d(\cdot)$ be the decoding neural network and assume

$$\begin{split} p(\mathbf{x}|\mathbf{z}) &= p(\mathbf{x}|\mathbf{z}^N) = \mathcal{N}(d(\mathbf{z}^N), \mathbf{I})),\\ \mathbf{d}(\mathbf{x}) &= d(\mathbf{z}^N), \quad \mathbf{z}^l = e^l(\mathbf{z}^{< l}, \mathbf{x}), \end{split}$$

the reconstruction term in ELBO_x becomes

$$\mathbb{E}_{q(\mathbf{z}|\mathbf{x})}\log p(\mathbf{x}|\mathbf{z}) = -\frac{1}{2} \|\mathbf{d}(\mathbf{x}) - \mathbf{x}\|^2.$$
(18)

Similarly, the reconstruction term in $ELBO_y$ is

$$\mathbb{E}_{q(\mathbf{z}|\mathbf{y})q(\mathbf{z}_{\mathbf{n}}|\mathbf{y})}\log p(\mathbf{y}|\mathbf{z}_{\mathbf{n}},\mathbf{z}) = -\frac{1}{2}\mathbb{E}_{q(\mathbf{z}_{\mathbf{n}}|\mathbf{y})}\|\mathbf{d}(\mathbf{y}) - \mathbf{y}\|^{2},$$
(19)

where we assume

$$\begin{split} p(\mathbf{y}|\mathbf{z}, \mathbf{z}_{\mathbf{n}}) &= p(\mathbf{y}|\mathbf{z}^{N}, \mathbf{z}_{\mathbf{n}}^{N}) = \mathcal{N}(d(\mathbf{z}^{N}, \mathbf{z}_{\mathbf{n}}^{N}), \mathbf{I})), \\ \mathbf{d}(\mathbf{y}) &= d(\mathbf{z}^{N}, \mathbf{z}_{\mathbf{n}}^{N}), \quad \mathbf{z}^{l} = e^{l}(\mathbf{z}^{< l}, \mathbf{y}), \\ \mathbf{z}_{\mathbf{n}}^{l} \sim \mathcal{N}(\mu_{q}^{l}(\mathbf{z}_{\mathbf{n}}^{< l}, \mathbf{y}), \sigma_{q}^{l}(\mathbf{z}_{\mathbf{n}}^{< l}, \mathbf{y})). \end{split}$$

The reparametrization technique [55] can be used for estimating (19). In this work, we set the number of layers N = 3. See Figure 3 for the details of the network architecture.

Remark 3. To synthetic the paired training data, it is sufficient to generate the degradation of clean images rather than the image content. Thus, we model the image part as an auto-encoder without generation ability and the degradation part as the hierarchical VAE model [60], [61]. This choice can strengthen the reconstruction ability of our model. A similar idea has been proposed in the VQ-VAE model [62], [63], which uses a discrete latent space and shows excellent generative ability. However, in the VQ-VAE model, it needs to construct a discrete codebook for p(z) and sample with a PixelCNN [64] model that increases the training difficulty. In our model, since there are no constraints on the inference model $q(z|\cdot)$ for the image part, we do not need to construct a codebook for latent variable z.

Inference invariant condition. The inference invariant condition (8) is a key tool for disentangling the image part and degradation part. In fact, it assumes that for the paired clean and noisy data $(\mathbf{x}, \mathbf{y}) \sim p(\mathbf{x}, \mathbf{y})$, the inferred latent variable \mathbf{z} is the same. This condition holds if we can extract the common information between \mathbf{x} and \mathbf{y} . In our method, we pass the input data through a pre-processing operator h, which can be manually designed or trained by an extra dataset. In LUD-VAE, we propose h based on the following principle. Given a paired data (\mathbf{x}, \mathbf{y}) and assume $\mathbf{y} = \mathbf{x} + \mathbf{n}$, we add Gaussian noise on both \mathbf{x} and \mathbf{y} :

$$h(\mathbf{x}) = \mathbf{x} + \mathbf{n}_{\mathbf{x}}, \quad h(\mathbf{y}) = \mathbf{x} + \mathbf{n} + \mathbf{n}_{\mathbf{y}},$$
 (20)

TABLE 1: Quantitative comparison on real-world super-resolution datasets AIM19 and NTIRE20. All synthetic datasets are trained with the ESRGAN model [11].

Mathad		AIM19		NTIRE20		
Methoa	PSNR ↑	SSIM \uparrow	LPIPS \downarrow	PSNR ↑	SSIM \uparrow	LPIPS \downarrow
ESRGAN-Bicubic	21.69	0.5517	0.517	20.45	0.3241	0.675
FSSR [30]	20.81	0.5242	0.387	21.07	0.4356	0.414
Impressionism [54]	21.99	0.6060	0.420	25.27	0.6731	0.229
DASR [31]	21.06	0.5658	0.375	23.70	0.5748	0.328
DeFlow-NP [1]	21.06	0.5842	0.346	24.81	0.6777	0.225
LUD-VAE (ours)	22.32	0.6197	0.341	25.72	0.7173	0.232
DeFlow [1]	22.25	0.6214	0.349	25.87	0.7005	0.218
CinCGAN [65]	21.60	0.6129	0.461	24.83	0.6752	0.509

Fig. 4: Three methods to generate paired data. C2N is used in LUD-VAE.

where $\mathbf{n_x} \sim \mathcal{N}(\mathbf{0}, \sigma_{\mathbf{x}}^2 \mathbf{I})$ and $\mathbf{n_y} \sim \mathcal{N}(\mathbf{0}, \sigma_{\mathbf{y}}^2 \mathbf{I})$. In this case, the inference model becomes

$$q(\mathbf{z}|\mathbf{x}) = \mathbb{E}_{p(\mathbf{n}_{\mathbf{x}})}q(\mathbf{z}|\mathbf{x} + \mathbf{n}_{\mathbf{x}}),$$

$$q(\mathbf{z}|\mathbf{y}) = \mathbb{E}_{p(\tilde{\mathbf{n}}_{\mathbf{y}})}q(\mathbf{z}|\mathbf{x} + \tilde{\mathbf{n}}_{\mathbf{y}}),$$
(21)

where $\tilde{\mathbf{n}}_{\mathbf{y}} = \mathbf{n} + \mathbf{n}_{\mathbf{y}} \sim \mathcal{N}(\mathbf{n}, \sigma_{\mathbf{y}}^2 \mathbf{I})$. Thus, the difference between $q(\mathbf{z}|\mathbf{x})$ and $q(\mathbf{z}|\mathbf{y})$ comes from the distance between $p(\mathbf{n}_{\mathbf{x}})$ and $p(\tilde{\mathbf{n}}_{\mathbf{y}})$, and it has

$$D_{\mathrm{KL}}(p(\tilde{\mathbf{n}}_{\mathbf{y}}) \| p(\mathbf{n}_{\mathbf{x}})) = \log \frac{\sigma_{\mathbf{x}}^{K}}{\sigma_{\mathbf{y}}^{K}} - \frac{K}{2} + \frac{K \sigma_{\mathbf{y}}^{2}}{2\sigma_{\mathbf{x}}^{2}} + \frac{\|\mathbf{n}\|_{2}^{2}}{2\sigma_{\mathbf{x}}^{2}}, \quad (22)$$

where *K* is the dimension of random variable $\mathbf{n}_{\mathbf{x}}$. The above KL divergence approaches to 0 as $\sigma_{\mathbf{x}}$, $\sigma_{\mathbf{y}}$ go to infinity. In this sense, it is reasonable that the inference invariant condition approximately hold. See the quantitative results Table 4 with different noise levels.

Remark 4. In this work, we use Monte Carlo method to estimate the above expectation terms in (21) and the number of sampling is set to 1 for each iterations.

Generate synthetic paired data. After the training procedure, there are three methods to generate paired data.

- 1) L2CN: Latent to clean and noisy pair. Sample $\mathbf{z_n} \sim p(\mathbf{z_n})$, $\mathbf{z} \sim p(\mathbf{z})$, then generate (\mathbf{x}, \mathbf{y}) from $p(\mathbf{y}|\mathbf{z}, \mathbf{z_n})$ and $p(\mathbf{x}|\mathbf{z})$.
- 2) N2C: Noisy to clean. Sample $\mathbf{y} \sim p(\mathbf{y})$, inference the latent variable \mathbf{z} with $q(\mathbf{z}|\mathbf{y})$, then generate the corresponding \mathbf{x} from $p(\mathbf{x}|\mathbf{z})$.
- 3) C2N: Clean to noisy. Sample $\mathbf{x} \sim p(\mathbf{x})$, inference the latent variable \mathbf{z} with $q(\mathbf{z}|\mathbf{x})$, sample $\mathbf{z}_{\mathbf{n}} \sim p(\mathbf{z}_{\mathbf{n}})$, then generate the corresponding \mathbf{y} from $p(\mathbf{y}|\mathbf{z}, \mathbf{z}_{\mathbf{n}})$.

See Figure 4 for the graphical explanations, and we adopt C2N to generate paired training data for the downstream tasks. This choice is consistent with the choice of the space of latent variable z since it does not have the generation ability. Moreover, the results in Table 3 show that C2N achieves better performance than the N2C method.

4 **EXPERIMENTS AND RESULTS**

We evaluate the performance of our LUD-VAE method on real-world image denoising and super-resolution tasks. First, we use LUD-VAE to learn the unknown degradation model under the unpaired learning settings, generate the synthetic training dataset, and then use an off-the-shelf supervised learning algorithm to learn the restoration model. All experiments are evaluated in the sRGB space.

4.1 Datasets and evaluation metrics

Two real-world super-resolution datasets and one realworld denoising dataset are chosen to evaluate our method: **AIM19**: Track 2 of the AIM 2019 real-world super-resolution challenge [2] provides a dataset of unpaired noisy-clean images. The noisy images are synthesized with an unknown combination of noise and compression, which mainly behave as structural and low frequency noises. The task is to learn a super-resolution model from the unpaired dataset, which restores high-resolution clean images from the lowresolution noisy inputs. The challenge also provides a validation set of 100 paired images, where different models can be compared with quantitative metrics. We refer to this dataset as the AIM19 dataset.

NTIRE20: Track 1 of NTIRE 2020 super-resolution challenge [33] follows the same setting as the AIM19 dataset, where it features an entirely different type of degradation, namely highly correlated high-frequency noise. As AIM19, there is a validation set containing 100 paired images that can be quantitatively evaluated. We refer to this dataset as the NTIRE20 dataset.

SIDD: The smartphone image denoising dataset (SIDD) [13] provides 30,000 noisy images from 10 scenes under different lighting conditions using five representative smartphone cameras and generates their ground truth images. We use the SIDD-Small Dataset and ignore the original index of clean noisy images to set up an unpaired dataset. It also provides the validation and benchmark datasets, each of which is cropped into 32 blocks of size 256×256 , resulting in 1024 image blocks in each dataset. We refer to this dataset as the SIDD dataset.

For all three datasets, we report the peak signal-to-noise ratio (PSNR) and the structural similarity index (SSIM) [66]. For AIM19 and NTIRE20 datasets, we also compute the LPIPS [67] distance, which is based on the comparison between features of a neural network. Here we use a pre-trained AlexNet [68] model.



Fig. 5: Visual comparison on real-world super-resolution dataset AIM19.

4.2 Implementation details

We train all LUD-VAE models for 200k iterations with the Adam [69] optimizer. The initial learning rate is set to 1e-4 and is halved after 100k iterations. We use a batch size of 16, containing random crops of size 64×64 . Batches are sampled randomly so that images from each domain are drawn with the same possibility. Random flips and rotates are used as data augmentation. Besides, to avoid the posterior collapse, we apply the KL annealing method [70], and use the linear anneal scheme for $D_{\rm KL}(q(\mathbf{z_n}|\mathbf{y}_i)||p(\mathbf{z_n}))$ in the first 10k iteration.

Super-resolution case. For real-world super-resolution, the degradation process is

$$\mathbf{y} = \mathcal{D}(\mathbf{x}) + \mathbf{n},\tag{23}$$

where \mathcal{D} is an unknown downsample operator, and **n** is an unknown noise. To simplify this problem, we replace \mathcal{D} by the bicubic downsample \mathcal{B} and put the approximation error of \mathcal{B} and \mathcal{D} into **n**, *i.e.*,

$$\mathbf{y} = \mathcal{B}(\mathbf{x}) + \mathbf{n}',\tag{24}$$

where $\mathbf{n}' = \mathbf{n} + \mathcal{T}(\mathbf{x}) - \mathcal{B}(\mathbf{x})$. Then for the given unpaired high-resolution clean dataset $\{\mathbf{x}_i\}$ and low-resolution noisy dataset $\{\mathbf{y}_j\}$, we first bicubic downsample $\{\mathbf{x}_i\}$ to low-resolution clean dataset $\{\mathcal{B}(\mathbf{x}_i)\}$, and then use LUD-VAE to

learn the noise n' from { $\mathcal{B}(\mathbf{x}_i)$ } and { \mathbf{y}_j }. After the training process, we use LUD-VAE to transfer the low-resolution clean dataset { $\mathcal{B}(\mathbf{x}_i)$ } to the synthetic degraded dataset, forming a paired training set with the high-resolution clean dataset { \mathbf{x}_i }. For the AIM19 dataset, we normalize the noisy image to make it has the same channel-wise mean and standard deviation as the clean domain, and then de-normalize the synthetic noisy image to constitute the training dataset. In AIM19, we set the $\sigma_{\mathbf{x}} = 15$, $\sigma_{\mathbf{y}} = 10$ in (20) for AIM19 and $\sigma_{\mathbf{x}} = 10$, $\sigma_{\mathbf{y}} = 5$ for NTIRE20.

Denoising case. For the denoising case, the degradation process becomes

$$\mathbf{y} = \mathbf{x} + \mathbf{n}.\tag{25}$$

We use LUD-VAE to learn the noise **n** from unpaired clean $\{\mathbf{x}_i\}$ and noisy $\{\mathbf{y}_j\}$ datasets directly. After training, we transform the clean image into the corresponding noisy image to obtain the paired training set. In this case, we set $\sigma_{\mathbf{x}} = 40$, $\sigma_{\mathbf{y}} = 20$ in (20).

4.3 Results in real-world super-resolution

We compare LUD-VAE with four unpaired degradation modeling methods namely FSSR [30] the winner of the AIM 2019 real-world super-resolution challenge [2]; Impressionism [54] the winner of the NTIRE 2020 real-world super-

Fig. 6: Visual comparison on real-world super-resolution dataset NTIRE20.



Fig. 7: The synthesis of training data in super-resolution.

resolution challenge [33]; DASR [31] a recently proposed GAN-based method; DeFlow [1] a flow-based method. We use the settings without the pre-trained network to retrain the DeFlow model, denoted as DeFlow-NP. We use LUD-VAE and these four methods to learn the unknown degradation model, then downsample the high-resolution images and generate the low-resolution noisy images to obtain the paired training dataset. In addition, we use the downsampled data without degradation as the baseline, denoted as ESRGAN-Bicubic, see Figure 7. When we have the training dataset, we use the real-world super-resolution model ESRGAN [11] to obtain the final super-resolution results for all methods. We use the training code from Impressionism and train the ESRGAN model for 60k iterations and choose the final model with the best LPIPS score on the validation dataset after 50k, 60k iterations. We also compare LUD-VAE with the unsupervised super-resolution model CinCGAN [65].

The quantitative results are shown in Table 1. For the evaluation metrics, PSNR and SSIM focus on the restoration of the overall content of the image, while LPIPS pays more attention to the image details, so these two types of metrics are mutually exclusive from each other. Many methods perform well on only one metric type, whereas LUD-VAE performs uniformly well on these three metrics. For the AIM19 dataset, the degradation mainly behaves as low-frequency and structured noise, while the NTIRE20 dataset is primarily high-frequency noise. We found that LUD-VAE can learn these different degradation models without paired data. It is worth noting that the performance of ESRGAN-Bicubic, which only uses bicubic downsampling without any other degradation, is not satisfactory (Figure 5 and Figure 6). Thus, the results illustrate the necessity of



Fig. 8: Visual comparison on real-world image denoising dataset SIDD benchmark.

TABLE 2: Quantitative comparison on real-world image denoising dataset SIDD

Mathad	SIDD bei	nchmark	SIDD va	SIDD validation	
Method	$\text{PSNR} \uparrow$	$\text{SSIM} \uparrow$	$\text{PSNR} \uparrow$	$\text{SSIM} \uparrow$	
CBM3D [71]	25.65	0.685	25.65	0.475	
MCWNNM [72]	33.37	0.875	33.40	0.815	
NC [73]	31.26	0.826	31.31	0.725	
N2S [51]	29.56	0.808	30.72	0.787	
AWGN	32.12	0.868	32.06	0.809	
DeFlow-NP [1]	33.54	0.875	33.53	0.817	
LUD-VAE (ours)	34.82	0.926	34.91	0.892	
DeFlow [1]	33.81	0.897	33.82	0.846	
DnCNN [6]	36.54	0.927	36.83	0.870	

considering the noisy model in real-world super-resolution tasks.

Comparison with flow-based methods. Our model achieves comparable results with the DeFlow model. On the AIM19 dataset, our model outperforms DeFlow on the PSNR and LPIPS metrics, and on the NTIRE20 dataset, our model outperforms the DeFlow model on the SSIM metric. It is worth noting that the DeFLow model requires the use of a pre-trained model, which is trained on paired low-

resolution and high-resolution data, which provides strong prior knowledge. Our method performs better than the DeFlow-NP model that does not use the pre-trained model. At the same time, the parameters of the DeFlow model are much larger than our model, which will bring difficulties to practical use, see Table 6. From the visual results in Figure 5 and Figure 6, our result removes the noise and is smoother and clearer than the DeFlow-NP model.

Comparison with GAN-based methods. We find that our model is more robust and effective. This may be because GAN-based methods often need to fine-tune different loss functions, and the cycle-consistency constraint is too weak to attain the theoretical guarantee for image restoration tasks. In the LUD-VAE model, we do not use any heuristic loss function, such as perceptual loss or GAN-style loss, to further improve the quantitative score. The visual results in Figure 5 and Figure 6 show that the performance of the GAN-based method is unstable; some images are still noisy, while our approach is more robust.

4.4 Results in real-world image denoising

We compare LUD-VAE with the unsupervised denoising method CBM3D [71], MCWNNM [72], NC [73], the datasetbased denoising method N2S [51], the degradation modeling method DeFlow/DeFlow-NP [1], and the fully supervised method DnCNN [6]. We also set up a baseline method as additive white Gaussian noise degradation, denoted as AWGN. Since the SIDD dataset contains images with different noise levels, we synthetic the degraded images with different noise levels using different degradation methods. For AWGN, we randomly apply Gaussian noise with zero mean and standard deviation $\sigma \in [20, 120]$ to each image; for DeFLow/DeFlow-NP, we randomly apply synthetic noise with the noise level parameter $t \in [1, 4]$ to each image; for LUD-VAE we choose $\sigma_{\mathbf{x}} \in [30, 70]$ in (20) for each image randomly; For AWGN, LUD-VAE, and DeFlow/DeFlow-NP, we use the DnCNN [6] for the downstream denoising tasks for 100k iterations with an initial learning rate 1e - 4and halved in 50k iteration. We test each method on the validation set every 500 iterations, and choose the final model with the best PSNR to evaluate on the benchmark set.

The quantitative results are shown in Table 2. We find that LUD-VAE has achieved the best results except for the supervised method DnCNN. DeFlow does not perform very well on this dataset; the reason may be the incomplete loss function of the model (see Appendix) and the instability of the training process. The dataset-based denoising method N2S performs poorly because the noise does not satisfy the spatial uncorrelated assumption. The visual results in Figure 8 show that our results are closer to DnCNN's, with complete noise removal compared with the other methods.

4.5 Ablation study and discussion

Generation method. We compare different methods in Figure 4 to generate paired data in our model. For L2CN, we need to train an additional PixelCNN [64] to sample from the prior distribution $p(\mathbf{z})$, which is inconvenient for practical use; thus, we do not adopt this method. For N2C and C2N, we evaluate them on the SIDD dataset. The results are shown in Table 3, and we find that C2N achieves better performance than N2C. One possible reason is that N2C is a deterministic mapping, while C2N is stochastic. Since one clean image has many different degraded counterparts, C2N is more suitable for the downstream denoising model.

Inference invariant condition. We verify the inference invariant condition on the AIM19 dataset. Firstly, we evaluate our method with different pre-processing noise levels $\sigma_{\mathbf{x}}$ and $\sigma_{\mathbf{y}}$, and then use the Maximum Mean Discrepancy [74] (MMD) to measure the distance between $q(\mathbf{z}|\mathbf{x})$ and $q(\mathbf{z}|\mathbf{y})$ for validating the establishment of the Inference invariant condition (8). We use the validation set of AIM19 and central crop the images to size 128×128 . For each image, we draw 10 samples $\{\mathbf{z}^i | \mathbf{x}\}_{i=1}^{10}, \{\mathbf{z}^i | \mathbf{y}\}_{i=1}^{10}$ from $q(\mathbf{z} | \mathbf{x}), q(\mathbf{z} | \mathbf{y})$ respectively, and then calculate MMD metric between $\{\mathbf{z}^i | \mathbf{x}\}_{i=1}^{10}$ and $\{\mathbf{z}^i | \mathbf{y}\}_{i=1}^{10}$. We compute the average MMD metric on the whole validation set, and the results are shown in Table 4. From the table, we find that in case $\sigma_{\mathbf{x}} = 0$, $\sigma_{\mathbf{v}} = 0$, the MMD is the largest, which means the inference invariant condition does not hold, and the model cannot learn the unknown degradation process, resulting in poor performance. In addition, the performance of $\sigma_{y} = 10, 15$ is better than $\sigma_{\mathbf{y}} = 5$. One possible explanation is that when $\sigma_{\mathbf{y}}$ is relatively large, the Gaussian noise will overwhelm the original

TABLE 3: Different generation methods on SIDD dataset.

Method	SIDD ber	nchmark	SIDD validation	
	$\text{PSNR} \uparrow$	$\text{SSIM} \uparrow$	$\text{PSNR} \uparrow$	$\text{SSIM} \uparrow$
N2C	34.74	0.901	34.78	0.858
C2N	34.82	0.926	34.91	0.892

TABLE 4: Validate inference invariant condition on AIM19 dataset.

$\sigma_{\mathbf{x}}$	$\sigma_{\mathbf{y}}$	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	MMD
0	0	21.66	0.5349	0.519	7.64
10	5	21.87	0.5510	0.374	4.42
15	10	22.32	0.6197	0.341	2.34
20	10	22.47	0.6355	0.389	2.08
15	5	22.96	0.6492	0.402	4.66
20	15	22.20	0.6199	0.343	1.35

unknown noise, making the inference invariant condition more satisfied, which is shown as the lower MMD metrics. Meanwhile, comparing the results of $\sigma_x = 15$, $\sigma_y = 10$ with $\sigma_{\mathbf{x}} = 20, \, \sigma_{\mathbf{y}} = 10$, we find that increasing the gap between $\sigma_{\mathbf{x}}$ and $\sigma_{\mathbf{y}}$, which means the degradation level of \mathbf{y} is larger will make LPIPS worse and make PSNR and SSIM better. This may be because when the degradation level of y is chosen to exceed the real degradation level, our model will synthesize higher degraded images than the real degraded images. The denoising performance of the image restoration model will trend to make the image smoother and lose the details. Increasing the pre-processing noise level can decrease the MMD metric, making the inference invariant condition more established. However, injecting higher-level noise will impair more image information, reducing the method's representation ability. Therefore, we need to balance inference invariant condition with an appropriate level of pre-processing noise and the model representation ability in practical applications.

Stochastic latent layer. We investigate the structure of our hierarchical VAE model with different stochastic layers, see Table 5 for the results. From the table, we find that the hierarchical VAE model outperforms the single layer structure significantly, indicating the necessity of using the hierarchical structure. Moreover, there is little difference in the final performance when the latent depth is larger than three. To reduce the number of parameters, our model adopts the three-layer structure.

Synthetic degradation comparison. We compute Fréchet Inception Distance [75] (FID) score to compare the quality of generated degradation images between different methods. FID is a metric that calculates the distance between the real

TABLE 5: Different stochastic layers of LUD-VAE on AIM19 dataset.

Num of Layer	PSNR ↑	SSIM \uparrow	LPIPS \downarrow
1 Layer	21.24	0.5364	0.375
2 Layers	21.96	0.6139	0.344
3 Layers	22.32	0.6197	0.341
4 Layers	22.14	0.6209	0.346
5 Layers	22.07	0.6137	0.341



Fig. 9: Synthetic degradation comparison with FID score.

TABLE 6: Model par	cameters and GPU	J running time.
--------------------	------------------	-----------------

Method	#Parameters (Million)	Times (Seconds)
FSSR [30]	1.62	0.011
DASR [31]	1.70	0.011
DeFlow [1]	62.94	2.548
LUD-VAE (ours)	5.71	0.387

and the generated images, where a smaller score means better generation quality. We central crop the images to 128×128 from the validation set of AIM19 and NTIRE20 datasets and compare the FID scores of different degradation modeling methods between the real and synthetic degradation images on these cropped images, and the results are shown in Figure 9. We find that the FID score of our method is better than other methods, except slightly inferior to Impressionism [54] on the NTIRE20 dataset, whereas the score of Impressionism on AIM19 is far worse than our method. One possible reason is that Impressionism is a handcrafted method with poor generalization. We also find that the FID score cannot fully reflect the performance of the downstream image restoration algorithm; for instance, the DeFlow [1] model has a poor score on the AIM19 dataset, but the restoration result outperforms other GAN-based models.

Model parameter. We compare the model parameters and the GPU running time for generating paired data (the dimension of the input image is $256 \times 256 \times 3$) of different unpaired learning models, see Table 6 for the results. The table shows that DASR and FSSR have the lowest model parameters and running times, but their restoration performance is beneath DeFlow and LUD-VAE methods. Compared with DeFlow, our model has much fewer parameters and running time, which is convenient for practical uses.

5 CONCLUSION

This paper proposes LUD-VAE, a degradation modeling method using unpaired data based on variational inference. We disentangle the clean and corrupted data domains through a probabilistic graphical model, which enables the transformation from clean to corrupted data by estimating the joint probability density function of the two domains using unpaired data from each domain individually. Moreover, we establish the equivalency between paired and unpaired learning for LUD-VAE under the inference invariant condition, which provides the mathematical rationale for our approach. We use LUD-VAE to generate synthetic training datasets for downstream supervised learning methods and evaluate them on real-world denoising and super-resolution tasks. Experimental results show that our method achieves state-of-the-art results on real-world image datasets.

ACKNOWLEDGMENTS

This work was supported by the National Key R&D Program of China (No. 2021YFA1001300), National Natural Science Foundation of China (No.11901338), Tsinghua University Initiative Scientific Research Program.

APPENDIX DISCUSSIONS ON MAXIMUM LIKELIHHOOD



Fig. 10: Generative process in DeFlow model.

Considering the generative model in Figure 10 that is suggested by the DeFlow method, where $\mathbf{z}_{\mathbf{x}} \sim \mathcal{N}(0, I)$, $\mathbf{z}_{\mathbf{y}} = \mathbf{z}_{\mathbf{x}} + \mathbf{u}$, $\mathbf{u} \sim \mathcal{N}(\mu_u, \Sigma_u)$. In the DeFlow method, it maximizes the log-likelihood function of the two marginal densities:

$$\max_{\theta} \sum_{i} \log p(\mathbf{x}_{i}) + \sum_{j} \log p(\mathbf{y}_{j}).$$
 (26)

where

$$\log p(\mathbf{x}) = \log |\det Df_{\theta}(\mathbf{x})| + \log \mathcal{N} (f_{\theta}(\mathbf{x}); 0, I)$$

$$\log p(\mathbf{y}) = \log |\det Df_{\theta}(\mathbf{y})| + \log \mathcal{N} (f_{\theta}(\mathbf{y}); \mu_u, I + \Sigma_u)$$
(27)

by the change of variables formula, where f_{θ} is an invertible normalizing flow. In the next, we show that the objective function in (26) is incomplete for representing the loglikelihood log $p(\mathbf{x}, \mathbf{y})$. Define

$$F(\mathbf{x}, \mathbf{y}) = (f_{\theta}(\mathbf{x}), f_{\theta}(\mathbf{y})) = (\mathbf{z}_{\mathbf{x}}, \mathbf{z}_{\mathbf{y}}), \quad (28)$$

where $(\mathbf{z}_{\mathbf{x}}, \mathbf{z}_{\mathbf{y}}) \sim \mathcal{N}\left(\begin{bmatrix} 0\\ \mu_u \end{bmatrix}, \begin{bmatrix} I & I\\ I & I + \Sigma_u \end{bmatrix}\right)$ as $\mathbf{z}_{\mathbf{y}} = \mathbf{z}_{\mathbf{x}} + \mathbf{u}$. Using the change of variables formula, we have

$$p(\mathbf{x}, \mathbf{y}) = |\det DF(\mathbf{x}, \mathbf{y})| \mathcal{N}\left(F(\mathbf{x}, \mathbf{y}); \begin{bmatrix} 0\\ \mu_u \end{bmatrix}, \begin{bmatrix} I & I\\ I & I + \Sigma_u \end{bmatrix}\right)$$
(29)

Since

$$F\left(\begin{bmatrix}\mathbf{x}\\\mathbf{y}\end{bmatrix}\right) = \begin{bmatrix}f_{\theta}(\mathbf{x})\\f_{\theta}(\mathbf{y})\end{bmatrix} \Rightarrow DF\left(\begin{bmatrix}\mathbf{x}\\\mathbf{y}\end{bmatrix}\right) = \begin{bmatrix}Df_{\theta}(\mathbf{x}) & 0\\0 & Df_{\theta}(\mathbf{y})\end{bmatrix},$$
(30)

we have

$$p(\mathbf{x}, \mathbf{y}) = \frac{|\det Df_{\theta}(\mathbf{x})| |\det Df_{\theta}(\mathbf{y})|}{\sqrt{(2\pi)^{2n} \det \Sigma_{u}}} \times \exp\left\{-\frac{1}{2} \begin{bmatrix} f_{\theta}(\mathbf{x}) \\ f_{\theta}(\mathbf{y}) - \mu_{u} \end{bmatrix}^{T} \begin{bmatrix} I + \Sigma_{u}^{-1} & -\Sigma_{u}^{-1} \\ -\Sigma_{u}^{-1} & \Sigma_{u}^{-1} \end{bmatrix} \begin{bmatrix} f_{\theta}(\mathbf{x}) \\ f_{\theta}(\mathbf{y}) - \mu_{u} \end{bmatrix}\right\}$$
(31)

where *n* is the dimension of random variables $\mathbf{z}_{\mathbf{x}}$ and $\mathbf{z}_{\mathbf{y}}$. Then the log-likelihood function $\log p(\mathbf{x}, \mathbf{y})$ can be decomposed into

$$\log p(\mathbf{x}, \mathbf{y}) = \log \mathcal{N} \left(f_{\theta}(\mathbf{x}); 0, I \right) + \log \left| \det Df_{\theta}(\mathbf{x}) \right| + \log \mathcal{N} \left(f_{\theta}(\mathbf{y}); \mu_{u}, I + \Sigma_{u} \right) + \log \left| \det Df_{\theta}(\mathbf{y}) \right| + \log \mathcal{N} \left(f_{\theta}(\mathbf{x}); -\mu_{u}, \Sigma_{u} \right) + f_{\theta}^{T}(\mathbf{x}) \Sigma_{u}^{-1} f_{\theta}(\mathbf{y}) + \frac{1}{2} \log \left((2\pi)^{n} \det(\Sigma_{u}) \right) - \frac{1}{2} \mu_{u}^{T} \Sigma_{u}^{-1} \mu_{u}.$$
(32)

Using marginal distribution notation, the maximal likelihood estimation is

$$\max_{\theta} \sum_{i} \log p(\mathbf{x}_{i}, \mathbf{y}_{i})$$

$$= \max_{\theta} \sum_{i} \log p(\mathbf{x}_{i}) + \sum_{j} \log p(\mathbf{y}_{j})$$

$$+ \sum_{i} \log \mathcal{N} \left(f_{\theta}(\mathbf{x}_{i}); -\mu_{u}, \Sigma_{u} \right) + f_{\theta}^{T}(\mathbf{x}_{i}) \Sigma_{u}^{-1} f_{\theta}(\mathbf{y}_{i})$$

$$+ \frac{1}{2} \log \left((2\pi)^{n} \det(\Sigma_{u}) \right) - \frac{1}{2} \mu_{u}^{T} \Sigma_{u}^{-1} \mu_{u}.$$
(33)

The additional term $\sum_i f_{\theta}^T(\mathbf{x}_i) \Sigma_u^{-1} f_{\theta}(\mathbf{y}_i)$ in (33) requires the paired information. But, in the DeFlow model, it further introduces conditional marginal likelihood and its relationship with conditional likelihood is still unknown and deserves further exploration. Inspired by the above derivations, we propose our generative graph in which the two latent variables are independent that is relatively easy for constructing an approximation of the log-likelihood.

REFERENCES

- V. Wolf, A. Lugmayr, M. Danelljan, L. Van Gool, and R. Timofte, "Deflow: Learning complex image degradations from unpaired data with conditional flows," in *CVPR*, 2021, pp. 94–103.
- [2] A. Lugmayr, M. Danelljan, R. Timofte, M. Fritsche, S. Gu, K. Purohit, P. Kandula, M. Suin, A. Rajagoapalan, N. H. Joon *et al.*, "Aim 2019 challenge on real-world image super-resolution: Methods and results," in *ICCVW*. IEEE, 2019, pp. 3575–3583.
- [3] A. Buades, B. Coll, and J.-M. Morel, "A non-local algorithm for image denoising," in *CVPR*, vol. 2. IEEE, 2005, pp. 60–65.
 [4] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denois-
- [4] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-d transform-domain collaborative filtering," *IEEE Trans Image Process*, vol. 16, no. 8, pp. 2080–2095, 2007.
- [5] S. Gu, L. Zhang, W. Zuo, and X. Feng, "Weighted nuclear norm minimization with application to image denoising," in *CVPR*, 2014, pp. 2862–2869.
- [6] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising," *IEEE Trans Image Process*, vol. 26, no. 7, pp. 3142–3155, 2017.
- [7] K. Zhang, W. Zuo, and L. Zhang, "Ffdnet: Toward a fast and flexible solution for cnn-based image denoising," *IEEE Trans Image Process*, vol. 27, no. 9, pp. 4608–4622, 2018.
- [8] S. Guo, Z. Yan, K. Zhang, W. Zuo, and L. Zhang, "Toward convolutional blind denoising of real photographs," in CVPR, 2019, pp. 1712–1722.
- [9] C. Dong, C. C. Loy, K. He, and X. Tang, "Learning a deep convolutional network for image super-resolution," in ECCV. Springer, 2014, pp. 184–199.

- [10] C. Dong, C. C. Loy, and X. Tang, "Accelerating the superresolution convolutional neural network," in ECCV. Springer, 2016, pp. 391–407.
- [11] X. Wang, K. Yu, S. Wu, J. Gu, Y. Liu, C. Dong, Y. Qiao, and C. Change Loy, "Esrgan: Enhanced super-resolution generative adversarial networks," in ECCVW, 2018, pp. 0–0.
- [12] A. Lugmayr, M. Danelljan, L. Van Gool, and R. Timofte, "Srflow: Learning the super-resolution space with normalizing flow," in *ECCV*. Springer, 2020, pp. 715–732.
- [13] A. Abdelhamed, S. Lin, and M. S. Brown, "A high-quality denoising dataset for smartphone cameras," in CVPR, 2018, pp. 1692– 1700.
- [14] A. Castillo, M. Escobar, J. C. Pérez, A. Romero, R. Timofte, L. Van Gool, and P. Arbelaez, "Generalized real-world superresolution through adversarial robustness," in *CVPR*, 2021, pp. 1855–1865.
- [15] S. Mohan, Z. Kadkhodaie, E. P. Simoncelli, and C. Fernandez-Granda, "Robust and interpretable blind image denoising via biasfree convolutional neural networks," arXiv:1906.05478, 2019.
- [16] D. Zheng, S. H. Tan, X. Zhang, Z. Shi, K. Ma, and C. Bao, "An unsupervised deep learning approach for real-world image denoising," in *International Conference on Learning Representations*, 2020.
- [17] A. Abdelhamed, M. A. Brubaker, and M. S. Brown, "Noise flow: Noise modeling with conditional normalizing flows," in CVPR, 2019, pp. 3165–3173.
- [18] R. D. Gow, D. Renshaw, K. Findlater, L. Grant, S. J. McLeod, J. Hart, and R. L. Nicol, "A comprehensive tool for modeling cmos imagesensor-noise performance," *IEEE Trans Electron Devices*, vol. 54, no. 6, pp. 1321–1329, 2007.
- [19] S. W. Hasinoff, F. Durand, and W. T. Freeman, "Noise-optimal capture for high dynamic range photography," in CVPR. IEEE, 2010, pp. 553–560.
- [20] J. Chen, J. Chen, H. Chao, and M. Yang, "Image blind denoising with generative adversarial network based noise modeling," in *CVPR*, 2018, pp. 3155–3164.
- [21] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Deep image prior," in CVPR, 2018, pp. 9446–9454.
- [22] Y. Quan, M. Chen, T. Pang, and H. Ji, "Self2self with dropout: Learning self-supervised denoising from single image," in CVPR, 2020, pp. 1890–1898.
- [23] X. Wu, M. Liu, Y. Cao, D. Ren, and W. Zuo, "Unpaired learning of deep image denoising," in ECCV. Springer, 2020, pp. 352–368.
- [24] S. Laine, T. Karras, J. Lehtinen, and T. Aila, "High-quality selfsupervised deep image denoising," *NIPS*, vol. 32, pp. 6970–6980, 2019.
- [25] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," *NIPS*, vol. 27, 2014.
- [26] A. Bulat, J. Yang, and G. Tzimiropoulos, "To learn image superresolution, use a gan to learn how to do image degradation first," in ECCV, 2018, pp. 185–200.
- [27] A. Lugmayr, M. Danelljan, and R. Timofte, "Unsupervised learning for real-world super-resolution," in *ICCVW*. IEEE, 2019, pp. 3408–3416.
- [28] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-toimage translation using cycle-consistent adversarial networks," in *ICCV*, 2017, pp. 2223–2232.
- [29] S. Bell-Kligler, A. Shocher, and M. Irani, "Blind super-resolution kernel estimation using an internal-gan," arXiv:1909.06581, 2019.
- [30] M. Fritsche, S. Gu, and R. Timofte, "Frequency separation for realworld super-resolution," in ICCVW. IEEE, 2019, pp. 3599–3608.
- [31] Y. Wei, S. Gu, Y. Li, R. Timofte, L. Jin, and H. Song, "Unsupervised real-world image super resolution via domain-distance aware training," in CVPR, 2021, pp. 13 385–13 394.
- [32] C. Chu, A. Zhmoginov, and M. Sandler, "Cyclegan, a master of steganography," arXiv:1712.02950, 2017.
- [33] A. Lugmayr, M. Danelljan, and R. Timofte, "Ntire 2020 challenge on real-world image super-resolution: Methods and results," in *CVPRW*, 2020, pp. 494–495.
- [34] C. Chen, Z. Xiong, X. Tian, and F. Wu, "Deep boosting for image denoising," in ECCV, 2018, pp. 3–18.
- [35] X. Jia, S. Liu, X. Feng, and L. Zhang, "Focnet: A fractional optimal control network for image denoising," in CVPR, 2019, pp. 6054– 6063.

- [36] K. Zhang, Y. Li, W. Zuo, L. Zhang, L. Van Gool, and R. Timofte, "Plug-and-play image restoration with deep denoiser prior," IEEE Trans. Pattern Anal. Mach. Intell., 2021.
- [37] T. Plotz and S. Roth, "Benchmarking denoising algorithms with real photographs," in CVPR, 2017, pp. 1586-1595.
- [38] J. Anaya and A. Barbu, "Renoir-a dataset for real low-light image noise reduction," J Vis Commun Image Represent, vol. 51, pp. 144-154, 2018.
- [39] A. Abdelhamed, R. Timofte, and M. S. Brown, "Ntire 2019 challenge on real image denoising: Methods and results," in CVPRW, 2019, pp. 0-0.
- [40] J. Cai, S. Gu, R. Timofte, and L. Zhang, "Ntire 2019 challenge on real image super-resolution: Methods and results," in CVPRW, 2019, pp. 0–0.
- [41] J. Cai, H. Zeng, H. Yong, Z. Cao, and L. Zhang, "Toward realworld single image super-resolution: A new benchmark and a new model," in ICCV, 2019, pp. 3086-3095.
- [42] S. Anwar and N. Barnes, "Real image denoising with feature attention," in CVPR, 2019, pp. 3155-3164.
- [43] D.-W. Kim, J. Ryun Chung, and S.-W. Jung, "Grdn: Grouped residual dense network for real image denoising and gan-based real-world noise modeling," in CVPRW, 2019, pp. 0-0.
- [44] Y. Zhang, K. Li, K. Li, L. Wang, B. Zhong, and Y. Fu, "Image superresolution using very deep residual channel attention networks," in ECCV, 2018, pp. 286–301.
- [45] J. Liang, J. Cao, G. Sun, K. Zhang, L. Van Gool, and R. Timofte, 'Swinir: Image restoration using swin transformer," in CVPR, 2021, pp. 1833-1844.
- [46] S. W. Zamir, A. Arora, S. Khan, M. Hayat, F. S. Khan, and M.-H. Yang, "Restormer: Efficient transformer for high-resolution image restoration," arXiv:2111.09881, 2021.
- [47] Z. Wang, X. Cun, J. Bao, and J. Liu, "Uformer: A general u-shaped transformer for image restoration," arXiv:2106.03106, 2021.
- [48] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang et al., "Photorealistic single image super-resolution using a generative adversarial network," in CVPR, 2017, pp. 4681-4690.
- [49] J. Lehtinen, J. Munkberg, J. Hasselgren, S. Laine, T. Karras, M. Aittala, and T. Aila, "Noise2noise: Learning image restoration without clean data," arXiv:1803.04189, 2018.
- [50] A. Krull, T.-O. Buchholz, and F. Jug, "Noise2void-learning denoising from single noisy images," in CVPR, 2019, pp. 2129-2137.
- [51] J. Batson and L. Royer, "Noise2self: Blind denoising by self-supervision," in *ICML*. PMLR, 2019, pp. 524–533.
- [52] M. Arjovsky and L. Bottou, "Towards principled methods for training generative adversarial networks," *arXiv*:1701.04862, 2017. [53] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative
- adversarial networks," in ICML. PMLR, 2017, pp. 214-223.
- [54] X. Ji, Y. Cao, Y. Tai, C. Wang, J. Li, and F. Huang, "Real-world super-resolution via kernel estimation and noise injection," in CVPRW, 2020, pp. 466-467.
- [55] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," arXiv:1312.6114, 2013.
- [56] Y. Zhao and C. Chen, "Unpaired image-to-image translation via latent energy transport," in CVPR, 2021, pp. 16418-16427.
- [57] M. Prakash, A. Krull, and F. Jug, "Fully unsupervised diversity denoising with convolutional variational autoencoders," arXiv:2006.06072, 2020.
- [58] A. Dosovitskiy and T. Brox, "Generating images with perceptual similarity metrics based on deep networks," NIPS, vol. 29, pp. 658-666, 2016.
- [59] C. K. Sønderby, T. Raiko, L. Maaløe, S. K. Sønderby, and O. Winther, "Ladder variational autoencoders," NIPS, vol. 29, 2016.
- [60] A. Vahdat and J. Kautz, "Nvae: A deep hierarchical variational autoencoder," arXiv:2007.03898, 2020.
- [61] R. Child, "Very deep vaes generalize autoregressive models and can outperform them on images," arXiv:2011.10650, 2020.
- [62] A. v. d. Oord, O. Vinyals, and K. Kavukcuoglu, "Neural discrete representation learning," arXiv:1711.00937, 2017.
- [63] A. Razavi, A. van den Oord, and O. Vinyals, "Generating diverse high-fidelity images with vq-vae-2," in NIPS, 2019, pp. 14866-14876
- [64] A. v. d. Oord, N. Kalchbrenner, O. Vinyals, L. Espeholt, A. Graves, and K. Kavukcuoglu, "Conditional image generation with pixelcnn decoders," arXiv:1606.05328, 2016.

- [65] Y. Yuan, S. Liu, J. Zhang, Y. Zhang, C. Dong, and L. Lin, "Unsupervised image super-resolution using cycle-in-cycle generative adversarial networks," in CVPRW, 2018, pp. 701-710.
- [66] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity. IEEE Trans Image Process, vol. 13, no. 4, pp. 600-612, 2004
- [67] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in CVPR, 2018, pp. 586-595.
- [68] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," NIPS, vol. 25, pp. 1097–1105, 2012.
- [69] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," arXiv:1412.6980, 2014.
- [70] S. R. Bowman, L. Vilnis, O. Vinyals, A. M. Dai, R. Jozefowicz, and S. Bengio, "Generating sentences from a continuous space," arXiv:1511.06349, 2015.
- [71] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Color image denoising via sparse 3d collaborative filtering with grouping constraint in luminance-chrominance space," in ICIP, vol. 1. IEEE, 2007, pp. I-313.
- [72] J. Xu, L. Zhang, D. Zhang, and X. Feng, "Multi-channel weighted nuclear norm minimization for real color image denoising," in ICCV, 2017, pp. 1096-1104.
- [73] M. Lebrun, M. Colom, and J.-M. Morel, "The noise clinic: a blind image denoising algorithm," Image Process. Line, vol. 5, pp. 1-54, 2015.
- [74] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola, "A kernel two-sample test," J Mach Learn Res, vol. 13, no. 1, pp. 723–773, 2012.
- M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochre-[75] iter, "Gans trained by a two time-scale update rule converge to a local nash equilibrium," NIPS, vol. 30, 2017.



Dihan Zheng is a Ph.D. student in in Yau Mathematical Sciences Center (YMSC), Tsinghua University supervised by Chenglong Bao. His research interests includes computer vision, image processing, generative modeling.



Xiaowen Zhang is a senior engineer in hisilicon. She received her Ph.D. from information and communication engineering, Shanghai Jiao Tong University in 2019. Her main research interests include computer vision, image processing.



Kaisheng Ma Kaisheng Ma is a tenure-track assistant professor of computer science in Institute for Interdisciplinary Information Sciences (IIIS), Tsinghua University. He leads the Algorithms, Architecture and Chipset Lab (ArChip). His research interests lie in the interdisciplinary fields of human-machine interface (HMI), brain-spired AI algorithm design, computer vision for selfdriving, compact model design, computer architecture and microelectronics. He believes in vertical integration for systems, with a special em-

phasis on implanted medical and neural devices, and chip level system solutions for self-driving systems. He received his Ph.D. in Computer Science and Engineering at the Pennsylvania State University, following Professor Vijaykrishnan Nanrayanan @ Penn State, Yuan Xie @ UCSB, Jack Sampson @ Penn State. His previous research on Nonvolatile Processor and Energy Harvesting won 2018 EDAA Best Dissertation Award. He publishes papers on conferences including NeuralPS, ICCV, AAAI, CVPR, ISCA, ASPLOS, MICRO, HPCA, DAC etc. He has won many awards, including: 2015 HPCA Best Paper Award, 2016 IEEE MICRO Top Picks, 2017 ASP-DAC Best Paper Award, 2016 IEEE MICRO Top Picks, 2017 ASP-DAC Best Paper Award, 2016 Penn State CSE Department Best Graduate Research, 2016 Cover Feature of NSF ASSIST Engineering Research Center, 2011 Yang Fuqing & Wang Yangyuan Academician Scholarship.



Chenglong Bao is an assistant professor in Yau mathematical sciences center, Tsinghua University and Yanqi Lake Beijing Institute of Mathematical Sciences and Applications. He received his Ph.D. from department of mathematics, National University of Singapore in 2014. His main research interests include mathematical image processing, large scale optimization and its applications.