

TAPE: Task-Agnostic Prior Embedding for Image Restoration

Lin Liu¹ Lingxi Xie³ Xiaopeng Zhang³ Shanxin Yuan⁴
Xiangyu Chen^{5,6} Wengang Zhou^{1,2} Houqiang Li^{1,2} Qi Tian³

¹CAS Key Laboratory of Technology in GIPAS,

EEIS Department, University of Science and Technology of China

²Institute of Artificial Intelligence, Hefei Comprehensive National Science Center

³Huawei Cloud BU

⁴Huawei Noah's Ark Lab

⁵University of Macau

⁶Shenzhen Institutes of Advanced Technology, CAS

Abstract. Learning a generalized prior for natural image restoration is an important yet challenging task. Early methods mostly involved hand-crafted priors including normalized sparsity, ℓ_0 gradients, dark channel priors, etc. Recently, deep neural networks have been used to learn various image priors but do not guarantee to generalize. In this paper, we propose a novel approach that embeds a task-agnostic prior into a transformer. Our approach, named Task-Agnostic Prior Embedding (TAPE), consists of two stages, namely, task-agnostic pre-training and task-specific fine-tuning, where the first stage embeds prior knowledge about natural images into the transformer and the second stage extracts the knowledge to assist downstream image restoration. Experiments on various types of degradation validate the effectiveness of TAPE. The image restoration performance in terms of PSNR is improved by as much as 1.45dB and even outperforms task-specific algorithms. More importantly, TAPE shows the ability of disentangling generalized image priors from degraded images, which enjoys favorable transfer ability to unknown downstream tasks.

1 Introduction

A good image prior can help to distinguish many kinds of noises from original image contents and improve the quality of images. Learning an image prior is important and challenging for image restoration tasks. Early studies explore specific degradation priors to achieve good performances on some low-level vision tasks, such as image dehazing [27, 78], image deblurring [46, 32], and image deraining [39, 77]. However, most priors are hand-crafted and mainly based on limited observations. With the popularity of deep learning, data-driven image priors estimated by combining conventional degradation properties with deep neural networks have been explored [34, 23, 71, 31, 21, 45, 11]. But these networks capturing task-specific priors, do not guarantee to generalize to unseen tasks.

Recently, there are also efforts in learning complicated priors for low-level vision tasks [56, 22, 47, 10, 38, 66]. These methods can be roughly grouped into

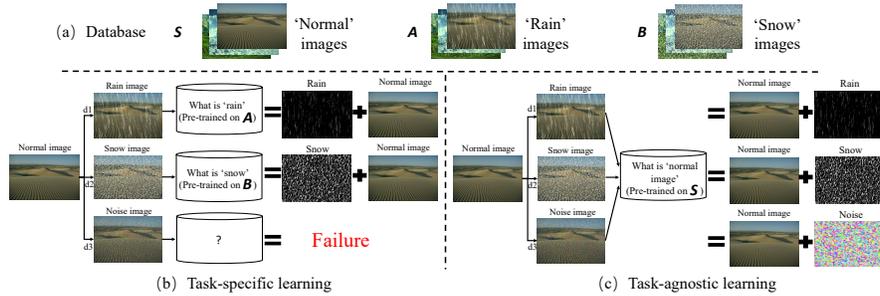


Fig. 1: The illustration of the differences of task specific learning and task-agnostic learning. Our method aims to learn ‘what is normal image’ instead of ‘what are rain, snow or other degradation’.

two types. The first type [10, 38] learns specific priors for each task, *i.e.*, the priors formulate ‘what is the distribution of specific noise’. Despite their effectiveness, these methods are often difficult to transfer across different tasks. The second type instead formulates generalized image priors, *i.e.*, ‘what is the distribution of normal images’. For this purpose, these methods [22, 47, 60] often make use of scalable GANs [6] pre-trained on natural images, hence, the learned priors are often hidden in a latent feature space, making it difficult to disentangle the noise from image contents, especially in the scenarios of complex images.

In this paper, we propose Task-Agnostic Prior Embedding (TAPE), a novel kind of priors that are easily generalized to different low-level vision tasks. An intuitive comparison between TAPE and prior task-specific learning is illustrated in Fig. 1. TAPE absorbs the benefits of the aforementioned approaches: on the one hand, we learn the distribution of normal images from non-degraded natural images, which does not rely on any true or synthesized degradation; on the other hand, the priors are encoded in a simple prior learning module named PLM and the main network can decode them by transformer decoders (query embeddings). The training procedure of TAPE consists of two stages, namely, task-agnostic pre-training and task-specific fine-tuning, where a pixel-wise contrastive loss is designed in the first stage for unsupervised low-level representation learning.

In the experiment, we pre-train our model on four tasks (including deraining, deraindrop, denoising, and demoireing), fine-tune and test it on these four known tasks and four unknown tasks (desnowing, shadow removal, super-resolution, and deblurring). After the one-time learning, the generalized image prior (through pre-training) can be transferred to different tasks. Quantitative and qualitative experimental comparisons show that the proposed TAPE improves the performance for multiple tasks in both task-specific and task-agnostic settings. In particular, our method improves the PSNR by 1.45dB, 1.03dB, 0.84dB, 0.49dB, and 0.75dB on the Rain200L, Rain200H, Raindrop800, SIDD, and TIP2018 datasets, respectively. The task-agnostic pre-training without touching the real noisy image on SIDD increases PSNR by 0.31dB. For the unseen tasks in the pre-training,

TAPe improves the PSNR by 0.91dB, 0.29dB, 0.41dB and 0.48dB on desnowing, shadow removal, super-resolution, and deblurring, respectively.

In summary, the contributions of our work are:

- The possibility and importance of learning task-agnostic and generalized image prior is addressed. As far as we know, TAPe is the first work to (explicitly) represent the universal prior that can be used in multiple image restoration tasks. We disentangle the generalized clean image prior of the corrupted images from the degrading objects/noises.
- We propose a two-stage method named TAPe for image restoration to learn the generalized degradation prior. The experiments demonstrate that our method can be easily applied to pre-train image restoration algorithms with other transformer backbones.
- We propose a pixel-wise contrastive loss in pre-training for learning better generalized features for PLM, which increases the generalization ability.

2 Related Work

Image Restoration. Image restoration is a general term for a series of low-level vision tasks, including denoising [70, 72, 24], deraining [39, 77, 37], deblurring [34, 30], demoireing [55, 75, 25], *etc.* The aim of image restoration is to restore clean \mathbf{x} from corrupted \mathbf{y} . The corrupted image \mathbf{y} can be formulated as, $\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{v}$, where \mathbf{H} , \mathbf{x} , and \mathbf{v} are degradation matrix, underlying clean image, and noise, respectively. Before the deep learning era, studies design hand-crafted features of the degradation objects (*e.g.*, rain, snow, *etc.*) for different image restoration tasks. With the popularity of convolutional neural networks (CNNs), a handful of deep-learning based methods are proposed to handle one or multiple types of image restoration tasks. Most of these methods design task-specific models or loss functions to achieve better performances. For image super-resolution [17, 28], Dong *et al.* propose SRCNN [17] to obtain high-resolution images from the corresponding low-resolution images. For HDR imaging, solutions [74, 15, 35] are proposed for using multiple exposed images to reconstruct an HDR image. Fu *et al.* [20] introduce a ResNet-based CNN for image deraining. Li *et al.* and Yu *et al.*, propose FDRNet [38] and RL-Restore [66] to handle hybrid-distorted image restoration tasks. Zheng *et al.* [76] propose a learnable bandpass filter network for image demoireing. Different from these methods, we explore the power of pre-training to handle several image processing tasks.

Image Degradation Prior and Natural Image Prior. Since image restoration is ill-posed, the image prior can help to constrain the solution space. From the Bayesian perspective, the solution $\hat{\mathbf{x}}$ can be obtained by optimizing:

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{y} - \mathbf{H}\mathbf{x}\|^2 + \lambda\Phi(\mathbf{x}), \quad (1)$$

where the first term is the fidelity and the second term is the regularization. Deep-learning based methods try to learn the prior parameters Θ and a compact

inference through an optimization of a loss function on a training dataset with corrupted-clean image pairs. Then Eqn. 1 can be refined as,

$$\min_{\Theta} \ell(\hat{\mathbf{x}}, \mathbf{x}) \quad \text{s.t.} \quad \hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{y} - \mathbf{H}\mathbf{x}\|^2 + \lambda\Phi(\mathbf{x}; \Theta). \quad (2)$$

Image priors have been widely used in computer vision, including markov random fields [53, 79], dark channel prior [27, 46], low rank prior [14], and total variation [54, 3]. He *et al.* [27] propose dark channel prior for image dehazing. It exploits the prior property that in an haze-free image there are pixels where at least one color channel is of low value. Chen *et al.* [14] propose a low-rank model to capture the spatial and temporal correlations between rain streaks. Different from these task-specific priors, we use a pre-trained network to extract more general priors from images. The prior queries are also being adjusted during end-to-end training. Recently, deep image prior (DIP) [56] shows that image statistics can be implicitly captured by the network’s structure, which is also a kind of prior. Inspired by DIP, some attempts use a pre-trained GAN as a source of image statistics [22, 47, 52, 19, 8, 60]. MGAN prior [22] utilizes multiple latent codes to increase the power of the pre-trained GAN model. DGP [47] fine-tunes the weights generator together with the latent code and use the discriminator to calculate the gap between the generated and real images.

Image Restoration Transformers. Transformer [57] is a new type of neural network framework of using mainly self-attention mechanism. It has achieved many successes in computer vision tasks, including object classification [18], objection detection [80, 16, 7], *etc.* Recently, it’s also been used in image restoration tasks [62, 10, 69, 40, 61, 36, 13, 42]. Chen *et al.* [10] and Li *et al.* [36] develop pre-trained transformers IPT and EDT respectively for some low-level vision tasks. Wang *et al.* [61] and Zamir *et al.* [67] design novel transformer structures, named Uformer and Restormer respectively, and obtain good performance on several image restoration tasks. Most of these methods proposed new transformer-based backbones for image restoration, whereas the goal of our work is a new pre-training pipeline for image restoration. One can apply our method to pre-train an image restoration algorithm with other backbones. We make use of the strong fitting ability of transformers and the learning ability of the transformer decoder to embed prior information.

3 Proposed Method

In this section, we first illustrate why we need to learn a task-agnostic prior (Sec. 3.1). And then, the architecture (Sec. 3.2) and the pipeline (Sec. 3.3) of our method are presented. At last, we briefly discuss the relationship to prior work (Sec. 3.4).

3.1 Motivation

The motivation is shown in Fig. 1, where we assume that many image restoration tasks (*e.g.*, deraining, desnowing, *etc.*) are present. Most existing image

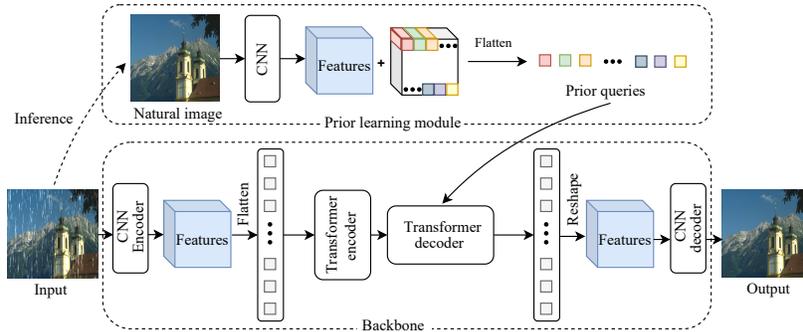


Fig. 2: The network architecture of our TAPE-Net. It consists of two parts: Backbone and prior learning module. With the input of natural images, PLM learns the features that natural images contain, not the features that noise contains. This makes our approach task-agnostic.

restoration methods learn a specific model for each single task, trying to capture task-specific priors (*e.g.*, what patterns are likely to be rain and thus need to be removed). While the learning procedure is straightforward, the trained models are difficult to be applied to new restoration tasks because the degradation patterns may have changed significantly.

To alleviate the burden, we propose a different pipeline that is generalized across different restoration tasks. The key is to learn **task-agnostic priors** (*i.e.*, what patterns are likely to be from clean, non-degraded images) rather than the aforementioned, task-specific counterparts. For this purpose, we embed an explicit module, the prior learning module (PLM), into the network, and design a two-stage learning procedure that (i) pre-trains the architecture on multi-source clean images and then (ii) fine-tunes it on specific image restoration datasets. The second stage often occupies a small portion of computation, showing the advantage of our method.

3.2 Network Architecture

The TAPE-Net (see Fig. 2) consists of two components: backbone and prior learning module. The backbone has a transformer architecture, containing a CNN encoder for feature extraction, a transformer encoder, a transformer decoder, and a CNN decoder for mapping the deep features into restored images. With the self-attention mechanism, the transformer can separate the generalized prior from the corrupted images. Different from conventional transformer [18], the decoder of our transformer takes additional prior queries, which comes from the prior learning module.

CNN encoder and CNN decoder. The CNN encoder consists of two 3×3 convolutional layers. The RGB image, $I \in \mathbb{R}^{3 \times H \times W}$, is the input of the CNN encoder, which generates a feature map $f_e \in \mathbb{R}^{64 \times H \times W}$ with 64 channels

and with the same resolution as I . The CNN decoder also consists of two 3×3 convolutional layers. It generates a reconstructed image $O \in \mathbb{R}^{3 \times H \times W}$.

Transformer encoder. The feature map f_e is firstly flattened into small patches $\{f_e^1, f_e^2, \dots, f_e^N\}$, where $f_e^i \in \mathbb{R}^{64P^2}$ ($i = 1, 2, \dots, N$), $N = \frac{HW}{P^2}$ is the total patch number and P is the patch size. A learnable position encoding PE_i with the same size of f_e^i is added to f_e^i and the sum (denoted as x_i) is sent into the transformer encoder. The transformer encoder has n transformer blocks ($n = 1$ in this work), each having a multi-head self-attention module and a feed forward network. The process of the transformer encoder can be formulated as,

$$\begin{aligned} x' &= \text{MSA}(\text{LN}(x), \text{LN}(x), \text{LN}(x)) + x \\ o_e &= \text{FFN}(\text{LN}(x')) + x', \end{aligned} \quad (3)$$

where MSA, FFN, and LN denote the multi-head self-attention module, feed forward network, and linear layer in the conventional transformer [57], respectively. $x = [x_1, x_2, \dots, x_N]$ and $o_e = [o_{e_1}, o_{e_2}, \dots, o_{e_N}]$ are the input and the output with the same size, respectively.

Prior learning module. The prior learning module (PLM) aims at providing additional prior queries to the transformer decoder. As shown in Fig. 2, PLM encodes an image into a feature map, and it can be formulated as $f_n = G_n(I_{gt})$, where f_n is a $64 \times H \times W$ feature map representing the deep natural image features and G_n is a VGG19 network to extract image features. Then f_n is flattened into a series of patches $[f_n^1, f_n^2, \dots, f_n^N]$ and combined with learnable parameters $[e_1, e_2, \dots, e_N]$ as follows,

$$Q = [e_1 + f_n^1, e_2 + f_n^2, \dots, e_N + f_n^N], \quad (4)$$

where Q have the same length as o_e .

Transformer decoder. The transformer decoder has a similar architecture as the transformer encoder except for an additional input of the prior queries Q (Similar with the object queries in [16]). In this paper, we use one transformer decoder block that consists of two multi-head self-attention (MSA) layers and one feed forward network (FFN). The transformer decoder is formulated as,

$$\begin{aligned} y &= \text{MSA}(\text{LN}(o_e) + Q, \text{LN}(o_e) + Q, \text{LN}(o_e)) + o_e \\ y' &= \text{MSA}(\text{LN}(y) + Q, \text{LN}(o_e), \text{LN}(o_e)) + y \\ o_d &= \text{FFN}(\text{LN}(y')) + y', \end{aligned} \quad (5)$$

where $o_d = [o_{d_1}, o_{d_2}, \dots, o_{d_N}]$ denotes the outputs of the transformer decoder. And then these patches are reshaped into f_d with the size of $64 \times H \times W$.

3.3 Optimization

3.3.1 The Training and Fine-tuning Pipeline

From Sec. 3.2, in our network design, PLM should extract statistics from the non-degraded (Ground truth) image and use the statistics to assist the main network for image restoration. However, ground truth is unavailable during the

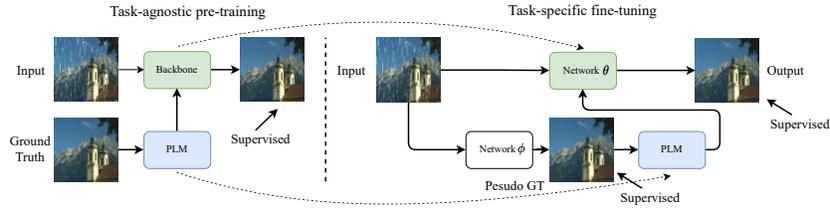


Fig. 3: The optimization procedure of TAPE. TAPE contains two stages: task-agnostic pre-training and task-specific fine-tuning. The dotted line means that the network trained in the previous stage is used to initialize the network in the next stage.

inference (test) stage. So PLM cannot extract statistics from it. To compensate, we train an auxiliary module (backbone ϕ in Fig. 3) to generate a pseudo GT from the degraded input. The pseudo GT, not being perfect, depicts the property of a non-degraded image to some extent. The pseudo GT is then fed into PLM for extracting general image priors, and the priors assist the main network (backbone θ) to generate the final output. As shown in Fig. 3, TAPE-Net has two stages: task-agnostic pre-training and task-specific fine-tuning.

In the task-agnostic pre-training, multiple low-level vision tasks are trained together, using corresponding datasets $\{D_1, \dots, D_m\}$, where D_i , ($i = 1, 2, \dots, m$) represents the dataset for task i . In each iteration, a pair of images (a corrupted image I_{cor} and its ground truth I_{gt}) are selected from one dataset D_i . The ground truth I_{gt} is sent into PLM to learn the prior queries, which are then sent to the backbone for end-to-end training. The combination of the L_1 loss and the proposed pixel-wise contrastive loss (see section 3.3.2) is used to optimize the network (the weighting parameter is for the latter). Due to the task-agnostic pre-training, both the backbone and PLM are well optimized.

As shown in the right part of Fig. 3, in the task-specific fine-tuning, because we cannot take the ground truth as input as discussed above, we use an auxiliary module (network ϕ in Fig. 3) to generate a pseudo GT from the degraded input (Note that in our paper, the network ϕ can be any neural network or the backbone borrowed from the pre-training stage. For faster convergence, we use the network pre-trained in the first stage). And then, the pseudo GT generated by the network ϕ is served as the input of pre-trained PLM. The PLM outputs the prior queries and help the pre-trained network θ to predict better final result. With the estimated pseudo ground truth, PLM can capture the natural image priors better. There are two ways to optimize the parameters in the fine-tuning stage, namely: 1) The network ϕ is fine-tuned by loss between pseudo GT and GT firstly, and then fixed when fine-tuning other networks; 2) All components are fine-tuned simultaneously. In the supplementary material, we will show that these two optimization methods lead to similar performance.

3.3.2 Pixel-wise Contrastive Loss

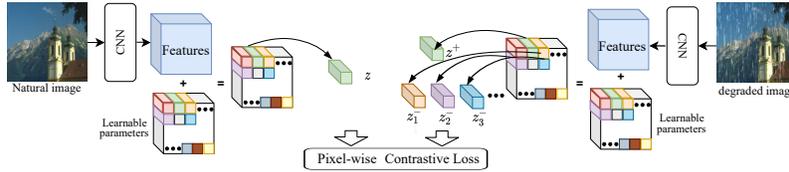


Fig. 4: The use of pixel-wise contrastive loss. z is selected from the prior queries of the natural image as ‘query’. And z^+ and z_j^- are selected as the ‘positive’ and ‘negative’ elements in the contrastive loss, respectively.

PLM aims to estimate the distribution of natural patches. However, due to the limited amount of training data, learning from the loss between prediction and GT (unary term) is insufficient for accurate estimation. Inspired by some self-supervised learning for high-level semantics (*e.g.* MoCo [26] and SimCLR [12]) and image to image translation method [4, 48], we propose a pixel-wise contrastive loss to offer another cue (binary terms) of estimation – the distance between the features of I_d and I_{gt} (from the same location) shall be smaller than that between features from different locations.

As shown in Fig. 4, in the task-agnostic pre-training stage, the degraded image I_d and the natural image I_{gt} are put into PLM, then Q^d and Q^{gt} are obtained as described in Sec. 3.2. We aim at minimizing the distance between the features of I_d and I_{gt} from the same location while maximizing the distance between features from different locations. For example, in Fig. 4, the roof without rain should be more closely associated with the roof contaminated by the rain than the other patches of the rainy input, such as other parts of the house or the blue sky.

Suppose that q_i^d is selected from $Q^d = \{q_1^d, q_2^d, \dots, q_N^d\}$ as the ‘query’ element in the contrastive loss. q_i^{gt} and $q_{j_1}^{gt}, q_{j_2}^{gt}, \dots, q_{j_m}^{gt}$ are selected from $Q^{gt} = \{q_1^{gt}, q_2^{gt}, \dots, q_N^{gt}\}$ as the ‘positive’ and ‘negative’ elements in the contrastive loss, respectively. Thus, the contrastive loss is formulated as,

$$\mathcal{L} = \sum_{t=1}^T \ell_t(q_i^d, q_i^{gt}, q_j^{gt}), \quad (6)$$

$$\ell(q_i^d, q_i^{gt}, q_j^{gt}) = -\log \left[\frac{\exp(q_i^d \cdot q_i^{gt} / \tau)}{\exp(q_i^d \cdot q_i^{gt} / \tau) + \sum_{k=1}^m \exp(q_i^d \cdot q_{j_k}^{gt} / \tau)} \right], \quad (7)$$

where $T = 256$ is the feature number we randomly choose each time and the temperature τ is set to 0.07. The negative sample number m is set to 256 in our work.

3.4 Relationship to Prior Work

1) Compared with recent multiple degradation prior learning methods (*e.g.*, IPT and EDT), our TAPE formulates generalized image priors, which means that our

Table 1: Datasets’ statistics (number of training and testing images) and quantitative comparison for two models (in terms of PSNR (dB)).

Dataset	Rain200L	Rain200H	Raindrop800	SIDD	TIP2018	Snow100K	ISTD	DIV2K	REDS
#Train/test images	1800/200	1800/200	800/60	96000/1280	10000/200	10000/500	1330/540	800/100	24000/3000
Baseline	31.72	22.81	26.85	37.41	26.77	25.42	26.28	31.25	32.46
TAPE-Net (Ours)	33.17	23.84	27.69	37.90	27.52	26.33	26.57	31.66	32.94
PSNR Gain	+1.45	+1.03	+0.84	+0.49	+0.75	+0.91	+0.29	+0.41	+0.48

method can generalize well to the pre-training-unknown tasks. Recently, AirNet [33] also has the ability to generalize to unknown tasks, but their learned representation contains the degraded information instead of normal image information through contrastive learning. 2) Different from the methods [22, 47, 5, 60] where the learned image priors are hidden in the parameters of the generator, the learned prior of our model is explicit. It makes easy for our method to disentangle the unwanted noise from the image contents in some complex image restoration cases.

4 Experiments and Analysis

In this section, we evaluate the performance of TAPE on several low-level vision tasks and conduct an ablation study.

4.1 Tasks and Datasets

For pre-training, we use five datasets, each for one type of degradation. We also test on four more datasets for four tasks that are unknown in the pre-training stage. For both training and testing, we resize images into the resolution of 256×256 , and then crop them into 64×64 patches for balancing the training procedure with different data sizes. Note that same resizing and cropping operations are also adopted for other models for a fair comparison. The evaluated tasks include denoising, deraining, deraindrop, demoireing, desnowing, shadow removal, super resolution and deblurring. As shown in Table 1 for details, the used datasets are: SIDD [1] for denoising, Rain200H and Rain200L [64] for deraining, Raindrop800 [49] for deraindrop, TIP2018 [55] for demoireing¹, Snow100K [43] for desnowing, ISTD [58] for shadow removal, DIV2K [2] for super resolution and REDS [44] for deblurring.

4.2 Implementation Details

Pre-training. We use one Nvidia Tesla V100 GPU to train our model using the Adam optimizer for 60×24000 iterations on the mixture of the five dataset

¹ We select a subset of TIP2018 and Snow100K with 10000 training image pairs and 200 test pairs; 10000 training image pairs and 500 test pairs, respectively.

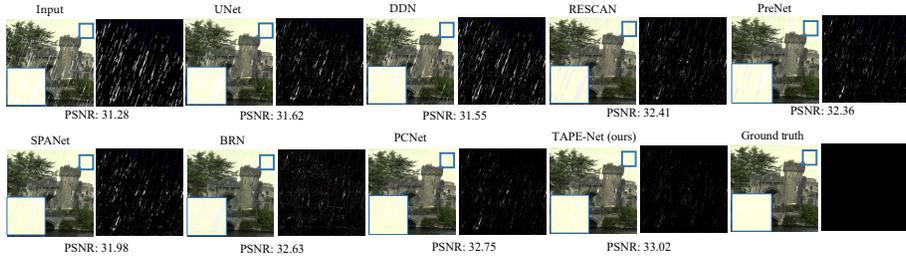


Fig. 5: Visual deraining comparison among our methods and SOTA deraining methods on Rain200L. The differences between the output and the ground truth are shown followed the predicted results.

(SIDD, Rain200L, Rain200H, Raindrop800, and TIP2018). The initial learning rate is set as 0.0002 and decayed to 0.0001 in the 20×24000 th iteration with batch size 4. In each iteration, we first randomly choose a dataset, from which one clean-corrupted image pair is randomly selected. **Fine-tuning.** After pre-training on all the datasets, we fine-tune the model on each desired task (*e.g.*, denoising). TAPE-Net is fine-tuned with 200 epochs and a learning rate of $2e-4$ for task-specific fine-tuning.

4.3 Pre-training & Generalization Ability

In this subsection, we illustrate that our method has good generalization performance on both pre-training-known tasks and pre-training-unknown tasks.

Pre-training-known tasks, corresponding data. To illustrate the effectiveness of our task-agnostic pre-training, we compare our pre-trained model with the model without pre-training (denoted as ‘Baseline’ in Table 1). The TAPE-Net improve the PSNR by 1.45dB, 1.03dB, 0.84dB, 0.49dB and 0.75dB on the Rain200L, Rain200H, Raindrop800, SIDD, and TIP2018 dataset, respectively. Please note that in image restoration, an 0.5dB improvement is usually considered significant. It demonstrates the effectiveness of the pre-training and the superiority of our model.

Pre-training-known tasks, different data. And we also do experiments to explore the good generalization performance of our pre-trained model when transferred to different distributions of data in the pre-training-known task. Our experiments show that pre-training with synthetic Gaussian noises helps to restore the images corrupted by real noises (see Table 3). In practice, in the task-agnostic pre-training, we use ground truth in SIDD with added synthetic Gaussian noises as input, without touching the real noise images on the SIDD dataset. Then the pre-trained model is fine-tuned on SIDD (using real-noise/non-noise image pairs). The PSNRs increase by 0.19dB, and 0.31dB respectively when the σ of the added Gaussian noise (Sampled from $\mathcal{N}(0, \sigma)$) are in $[5, 20]$, and $[1, 50]$, respectively. When the range of the added Gaussian noise is larger, the

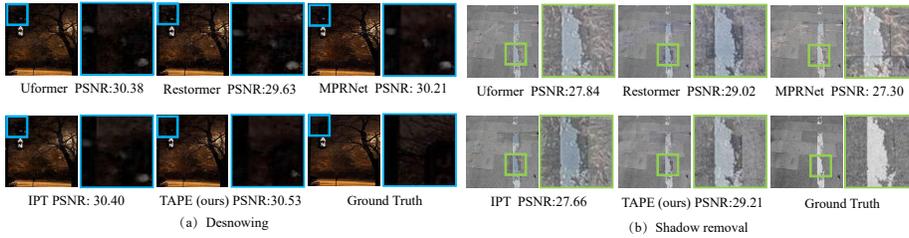


Fig. 6: Visual desnowing and shadow removal comparison among our methods and SOTA general image restoration methods on Snow100K and ISTD.

Table 2: Quantitative comparison with 4 SOTA general image restoration methods with similar model sizes (in terms of PSNR (dB)).

	Model size	Snow100K	ISTD	Rain200L	SIDD	Raindrop800
IPT	2.51M	26.26	26.34	32.67	38.80	27.86
Restormer	0.93M	26.80	26.42	33.61	38.91	27.98
UFormer	0.97M	26.50	26.27	32.66	38.84	27.70
MPRNet	1.11M	26.30	26.23	33.30	38.89	28.13
TAPE-Swin (Ours)	0.97M	26.93	26.61	34.46	38.98	29.15
TAPE-Restormer (Ours)	1.07M	26.91	26.65	34.28	39.01	29.18

generalization ability of the model in the pre-training stage is stronger, and the performance in the fine-tuning stage is better.

Pre-training-unknown tasks. To illustrate the generalization ability of our model, we conduct several experiments on the tasks that are unknown to the pre-training stage. In practice, we fine-tune the pre-trained model on four new low-level vision tasks: desnowing, shadow removal, super resolution and deblurring. As shown in Table 1, compared with the none-pre-trained model, TAPE-Net improves the PSNR by 0.91dB, 0.29dB, 0.41dB and 0.48dB, which demonstrates that the pre-trained model can capture more useful information and features from natural images. Learning task-agnostic priors and pixel-wise contrastive loss on pre-training stage can help the performance of fine-tuning on the unknown tasks.

4.4 Comparisons with State-of-the-Arts

In this subsection, the comparison between our methods and the very recent SOTA image restoration methods are shown in Table 2. we compare our methods (TAPE-swin and TAPE-restormer) with the 4 SOTA methods (IPT [10], Restormer [67], UFormer [61], MPRNet [68]) with the similar model parameters in 3 pre-training-known tasks and 2 pre-training-unknown tasks. The visual results on two pre-training-unknown tasks (desnowing and shadow removal) are shown in Fig. 6. Our method removes the artifacts more thoroughly and retains more details. In the supplementary material, we compare our TAPE-Net with 12

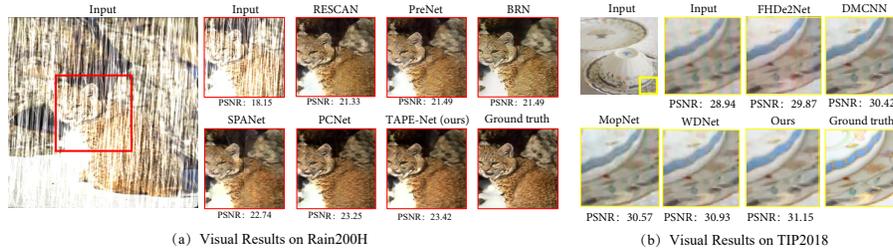


Fig. 7: Visual deraining comparison (a) and visual demoireing comparison (b) among our methods and task-specific methods.

Table 3: Quantitative comparison on SIDD. It shows the good generalization performance of TAPE-Net when transferred to different distributions of data in the pre-training-known task.

TAPE-Net W/o pre-train	Pre-train ($\sigma \in [5,20]$)	Pre-train ($\sigma \in [1,50]$)	
PSNR	37.41	37.60	37.72

Table 4: Ablation study of the impact of the multi-task pre-training, which illustrates that adding more tasks will not harm the performance.

	Raindrop800	Rain200L	TIP2018	Snow100K
RD+RL+RH	27.71	33.20	27.56	26.17
RD+TIP+RL+RH	27.70	33.18	27.54	26.26
RD+TIP+RL+RH+S	27.69	33.17	27.52	26.33

state-of-the-art task-specific methods, including deraining methods (DDN [20], SPANet [59], RESCAN [37], PreNet [51], BRN [50], SPDNet [65] and PCNet [29]); demoireing methods (DMCNN [55], MopNet [25], FHD2Net [25], HRDN [63], WDNNet [41] and MBCNN [76]); and denoising methods (DnCNN [70], FFDNet [72], RDN [73], and SADNet [9]) on PSNR. Qualitative results on deraining are shown in Figs. 5 and 7(a), showing that our methods get cleaner images and recover more details. The visual demoireing results are shown in Fig. 7(b). Ours can remove moire patterns successfully and restore the underlying clean image.

In order to show the performance of our method on pre-training-unknown tasks, we compare ours with the existing SOTA multi-task pre-training method, IPT [10]. We fine-tuned and tested the official pre-trained model of IPT on three unseen tasks, namely, desnowing, shadow removal and deblurring. We compared the gain of PSNR with and without pre-training, because the training patches and model size of IPT and ours are different. The pre-training stage of IPT boosts PSNR by 0.07dB, 0.19dB and 0.08dB, respectively, while the improvements of TAPE (our method) are 0.91dB, 0.29dB and 0.48dB, larger than that of IPT.

Table 5: The ablation study of the importance of each pre-trained part on the Rain200H and the Raindrop-TestB dataset. ‘-’ and ‘✓’ in the first three columns mean that the model parameters are randomized and pre-trained respectively before the finetuning. ‘✗’ means the corresponding part does not exist.

Name	Depth	network ϕ	network θ	PLM	Rain200H	Raindrop-TestB	Model size (M)
Baseline	3	✗	-	✗	25.57	26.13	0.76
With no-pre-trained all parts	2	-	-	-	25.46	25.96	1.12
With no-pre-trained ϕ	2	-	✓	✓	26.15	26.39	1.12
With no-pre-trained θ	2	✓	-	✓	25.96	26.35	1.12
With no-pre-trained θ and PLM	2	✓	-	-	25.73	26.17	1.12
Full model	2	✓	✓	✓	26.18	26.41	1.12

4.5 Ablation Study

Impact of the multi-task pre-training. We do ablation study to analyze the effect of the number of datasets in the pre-training. We pre-train our models on fewer datasets and compare with our original models. As shown in Table 4, RD, TIP, RL, RH, and S mean Raindrop800, TIP2018, Rain200L, Rain200H, and SIDD datasets respectively. ‘+’ means we use these datasets in the pre-training stage. We do the experiments on three pre-trained-known datasets and the maximum PSNR difference is 0.04 dB. This PSNR difference is within a controllable error range. Increasing the dataset in the pre-training is meaningful. Compared with pre-training with 3 datasets, pre-training with 5 datasets increases the PSNR by 0.16dB on Snow100K.

Importance of pre-training of each component. In order to verify which component is more important with pre-training (network ϕ , network θ or PLM), we randomize the weights of corresponding parts of TAPE-Net before the fine-tuning stage. As shown in Table 5, we do ablation study on Rain200H and Raindrop-TestB. All the modified models in this table have the same backbone, TAPE-Swin. Compared ‘With no-pre-trained ϕ ’ and the ‘Full model’, we can see that the PSNRs drop slightly (0.03dB and 0.02dB). With no-pre-trained θ effects more than the network ϕ (the PSNRs drop by 0.22dB and 0.06dB). ‘Without pre-training the network θ and PLM’ also makes the PSNRs drop a lot. Thus, the last four lines in Table 5 demonstrate that the effectiveness of pre-training of θ and PLM. The first two lines in Table 5 show that even if the model size is larger than Baseline, the performance is not good without pre-training.

The ablation studies about the pixel-wise contrastive loss, optimization methods in task-specific fine-tuning can be see in the supplementary material.

4.6 Visualization Results

To validate that PLM learns useful and meaningful features with our proposed pipeline, we visualize the features learned by PLM on the deraining task. We put the input, pseudo GT and GT into PLM to get their respective output Q .

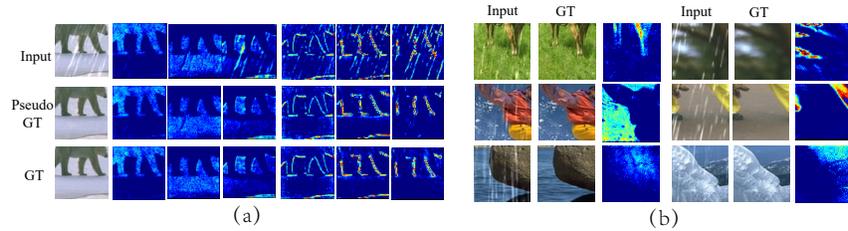


Fig. 8: The visualization of the PLM’s output, Q . (a) The features of the pseudo GT and GT are very similar, but most of the features of input have the features of degraded objects. It means that the features of the pseudo GT is useful and can help the image restoration of the backbone. (b) The last lines are one of the predicted results of PLM.

As shown in Fig. 8(a), in the output features, some channels tend to preserve the information of shapes (the first three feature maps) and edges (the last three feature maps). The features of the pseudo GT and GT are very similar, but most of the features of input have the features of degraded objects. It means that the features of the pseudo GT is useful and can help the image restoration of the backbone. From Fig. 8(b), we can see that with the help of pre-training, the PLM module can correlate the information of similar textures or patches from a long distance. Thus, the transformer decoder of the backbone can utilize these long-distance similar areas/patches to restore the image.

5 Conclusions and Limitations

Conclusions. In this paper, we address the possibility and importance of learning task-agnostic and generalized image prior. We propose a pipeline named TAPE to learn task-agnostic prior embedding for image restoration. TAPE has two stages: task-agnostic pre-training and task-specific fine-tuning. Our task-agnostic training strategy is able to learn generalized natural image prior. It has good generalization performance when faced with pre-training-unknown tasks.

Limitations. Although our method shows generalized ability on a few image restoration tasks, the learned statistics in PLM are still difficult to explain, either in theory or by visualization results. Without these proofs, the PSNR and SSIM numbers are only side evidences of the effectiveness of the task-agnostic priors. In the future, we will continue exploring the possibility of disentangling the task-agnostic priors as well as finding more essential ways to evaluate cross-task image restoration. The pipeline may make the training time longer than the baseline. And the performance of TAPE on mixed degradation tasks needs to be explored.

Acknowledgements. This work was supported by the National Natural Science Foundation of China under Contract 61836011 and 62021001. It was also supported by the GPU cluster built by MCC Lab of Information Science and Technology Institution, USTC.

References

1. Abdelhamed, A., Lin, S., Brown, M.S.: A high-quality denoising dataset for smart-phone cameras. In: CVPR, 2018
2. Agustsson, E., Timofte, R.: Ntire 2017 challenge on single image super-resolution: Dataset and study. In: CVPRW, 2017
3. Babacan, S.D., Molina, R., Katsaggelos, A.K.: Variational bayesian blind deconvolution using a total variation prior. TIP, 2008
4. Baek, K., Choi, Y., Uh, Y., Yoo, J., Shim, H.: Rethinking the truly unsupervised image-to-image translation. In: International Conference on Computer Vision (ICCV, 2021)
5. Bau, D., Strobel, H., Peebles, W., Zhou, B., Zhu, J.Y., Torralba, A., et al.: Semantic photo manipulation with a generative image prior. arXiv preprint arXiv:2005.07727 (2020)
6. Brock, A., Donahue, J., Simonyan, K.: Large scale gan training for high fidelity natural image synthesis. In: ICLR, 2018
7. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: ECCV, 2020
8. Chan, K.C., Wang, X., Xu, X., Gu, J., Loy, C.C.: Glean: Generative latent bank for large-factor image super-resolution. arXiv preprint arXiv:2012.00739 (2020)
9. Chang, M., Li, Q., Feng, H., Xu, Z.: Spatial-adaptive network for single image denoising. ECCV, 2020
10. Chen, H., Wang, Y., Guo, T., Xu, C., Deng, Y., Liu, Z., Ma, S., Xu, C., Xu, C., Gao, W.: Pre-trained image processing transformer. CVPR, 2021
11. Chen, L., Fang, F., Wang, T., Zhang, G.: Blind image deblurring with local maximum gradient prior. In: CVPR, 2019
12. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.E.: A simple framework for contrastive learning of visual representations. In: ICML, 2020
13. Chen, X., Wang, X., Zhou, J., Dong, C.: Activating more pixels in image super-resolution transformer. arXiv preprint arXiv:2205.04437 (2022)
14. Chen, Y.L., Hsu, C.T.: A generalized low-rank appearance model for spatio-temporally correlated rain streaks. In: ICCV, 2013
15. Dai, T., Li, W., Cao, X., Liu, J., Jia, X., Leonardis, A., Yan, Y., Yuan, S.: Wavelet-based network for high dynamic range imaging. arXiv preprint 2108.01434 (2021)
16. Dai, Z., Cai, B., Lin, Y., Chen, J.: Up-detr: Unsupervised pre-training for object detection with transformers. In: CVPR, 2021
17. Dong, C., Loy, C.C., He, K., Tang, X.: Image super-resolution using deep convolutional networks. TPAMI, 2015
18. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929, 2020
19. El Helou, M., Süsstrunk, S.: BIGPrior: Towards decoupling learned prior hallucination and data fidelity in image restoration. arXiv preprint arXiv:2011.01406 (2020)
20. Fu, X., Huang, J., Zeng, D., Huang, Y., Ding, X., Paisley, J.: Removing rain from single images via a deep detail network. In: CVPR, 2017
21. Golts, A., Freedman, D., Elad, M.: Unsupervised single image dehazing using dark channel prior loss. TIP, 2020

22. Gu, J., Shen, Y., Zhou, B.: Image processing using multi-code gan prior. In: CVPR, 2020
23. Guo, S., Liang, Z., Zhang, L.: Joint denoising and demosaicking with green channel prior for real-world burst images. arXiv preprint arXiv:2101.09870 (2021)
24. Guo, S., Yan, Z., Zhang, K., Zuo, W., Zhang, L.: Toward convolutional blind denoising of real photographs. In: CVPR, 2019
25. He, B., Wang, C., Shi, B., Duan, L.Y.: Mop moire patterns using mopnet. In: ICCV, 2019
26. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: CVPR, 2020
27. He, K., Sun, J., Tang, X.: Single image haze removal using dark channel prior. TPAMI, 2010
28. Isobe, T., Li, S., Jia, X., Yuan, S., Slabaugh, G., Xu, C., Li, Y.L., Wang, S., Tian, Q.: Video super-resolution with temporal group attention. In: CVPR, 2020
29. Jiang, K., Wang, Z., Yi, P., Chen, C., Lin, C.W.: Pcnnet: Progressive coupled network for real-time image deraining. In: TIP, 2021
30. Kupyn, O., Martyniuk, T., Wu, J., Wang, Z.: Deblurgan-v2: Deblurring (orders-of-magnitude) faster and better. In: ICCV, 2019
31. Lee, H., Sohn, K., Min, D.: Unsupervised low-light image enhancement using bright channel prior. IEEE Signal Processing Letters, 2020
32. Levin, A., Weiss, Y., Durand, F., Freeman, W.T.: Understanding and evaluating blind deconvolution algorithms. In: CVPR, 2009
33. Li, B., Liu, X., Hu, P., Wu, Z., Lv, J., Peng, X.: All-in-one image restoration for unknown corruption. In: CVPR (2022)
34. Li, L., Pan, J., Lai, W.S., Gao, C., Sang, N., Yang, M.H.: Blind image deblurring via deep discriminative priors. IJCV, 2019
35. Li, W., Xiao, S., Dai, T., Yuan, S., Wang, T., Li, C., Song, F.: S_j-hd²r: Selective joint high dynamic range and denoising imaging for dynamic scenes. arXiv preprint 2206.09611 (2022)
36. Li, W., Lu, X., Lu, J., Zhang, X., Jia, J.: On efficient transformer and image pre-training for low-level vision. In: arXiv preprint 2112.10175
37. Li, X., Wu, J., Lin, Z., Liu, H., Zha, H.: Recurrent squeeze-and-excitation context aggregation net for single image deraining. In: ECCV, 2018
38. Li, X., Jin, X., Lin, J., Liu, S., Wu, Y., Yu, T., Zhou, W., Chen, Z.: Learning disentangled feature representation for hybrid-distorted image restoration. In: ECCV, 2020
39. Li, Y., Tan, R.T., Guo, X., Lu, J., Brown, M.S.: Rain streak removal using layer priors. In: CVPR, 2016
40. Liang, J., Cao, J., Sun, G., Zhang, K., Van Gool, L., Timofte, R.: Swinir: Image restoration using swin transformer. In: ICCVW, 2021
41. Liu, L., Liu, J., Yuan, S., Slabaugh, G., Leonardis, A., Zhou, W., Tian, Q.: Wavelet-based dual-branch network for image demoiréing. ECCV, 2020
42. Liu, L., Yuan, S., Liu, J., Guo, X., Yan, Y., Tian, Q.: Siamtrans: Zero-shot multi-frame image restoration with pre-trained siamese transformers. In: AAAI, 2022
43. Liu, Y.F., Jaw, D.W., Huang, S.C., Hwang, J.N.: Desnownet: Context-aware deep network for snow removal. TIP, 2018
44. Nah, S., Baik, S., Hong, S., Moon, G., Son, S., Timofte, R., Lee, K.M.: Ntire 2019 challenge on video deblurring and super-resolution: Dataset and study. In: CVPRW, 2019
45. Pan, J., Bai, H., Tang, J.: Cascaded deep video deblurring using temporal sharpness prior. In: CVPR, 2020

46. Pan, J., Sun, D., Pfister, H., Yang, M.H.: Blind image deblurring using dark channel prior. In: CVPR, 2016
47. Pan, X., Zhan, X., Dai, B., Lin, D., Loy, C.C., Luo, P.: Exploiting deep generative prior for versatile image restoration and manipulation. In: ECCV, 2020
48. Park, T., Efros, A.A., Zhang, R., Zhu, J.Y.: Contrastive learning for unpaired image-to-image translation. In: ECCV, 2020
49. Qian, R., Tan, R.T., Yang, W., Su, J., Liu, J.: Attentive generative adversarial network for raindrop removal from a single image. In: CVPR, 2018
50. Ren, D., Shang, W., Zhu, P., Hu, Q., Meng, D., Zuo, W.: Single image deraining using bilateral recurrent network. TIP, 2020
51. Ren, D., Zuo, W., Hu, Q., Zhu, P., Meng, D.: Progressive image deraining networks: A better and simpler baseline. In: CVPR, 2019
52. Richardson, E., Alaluf, Y., Patashnik, O., Nitzan, Y., Azar, Y., Shapiro, S., Cohen-Or, D.: Encoding in style: a stylegan encoder for image-to-image translation. arXiv preprint arXiv:2008.00951 (2020)
53. Roth, S., Black, M.J.: Fields of experts: A framework for learning image priors. In: CVPR, 2005
54. Rudin, L.I., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. *Physica D: nonlinear phenomena*, 1992
55. Sun, Y., Yu, Y., Wang, W.: Moiré photo restoration using multiresolution convolutional neural networks. TIP, 2018
56. Ulyanov, D., Vedaldi, A., Lempitsky, V.: Deep image prior. In: CVPR, 2018
57. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. arXiv preprint arXiv:1706.03762, 2017
58. Wang, J., Li, X., Yang, J.: Stacked conditional generative adversarial networks for jointly learning shadow detection and shadow removal. CVPR, 2018
59. Wang, T., Yang, X., Xu, K., Chen, S., Zhang, Q., Lau, R.W.: Spatial attentive single-image deraining with a high quality real rain dataset. In: CVPR, 2019
60. Wang, X., Li, Y., Zhang, H., Shan, Y.: Towards real-world blind face restoration with generative facial prior. In: CVPR, 2021
61. Wang, Z., Cun, X., Bao, J., Liu, J.: Uformer: A general u-shaped transformer for image restoration. arXiv preprint 2106.03106
62. Yang, F., Yang, H., Fu, J., Lu, H., Guo, B.: Learning texture transformer network for image super-resolution. In: CVPR, 2020
63. Yang, S., Lei, Y., Xiong, S., Wang, W.: High resolution demoire network. In: ICIP, 2020
64. Yang, W., Tan, R.T., Feng, J., Liu, J., Guo, Z., Yan, S.: Deep joint rain detection and removal from a single image. In: CVPR, 2017
65. Yi, Q., Li, J., Dai, Q., Fang, F., Zhang, G., Zeng, T.: Structure-preserving deraining with residue channel prior guidance. ICCV (2021)
66. Yu, K., Dong, C., Lin, L., Loy, C.C.: Crafting a toolchain for image restoration by deep reinforcement learning. In: CVPR, 2018
67. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H.: Restormer: Efficient transformer for high-resolution image restoration. In: CVPR, 2022
68. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H., Shao, L.: Multi-stage progressive image restoration. In: CVPR (2021)
69. Zeng, Y., Fu, J., Chao, H.: Learning joint spatial-temporal transformations for video inpainting. In: ECCV, 2020
70. Zhang, K., Zuo, W., Chen, Y., Meng, D., Zhang, L.: Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. TIP, 2017

71. Zhang, K., Zuo, W., Gu, S., Zhang, L.: Learning deep cnn denoiser prior for image restoration. In: CVPR, 2017
72. Zhang, K., Zuo, W., Zhang, L.: Ffdnet: Toward a fast and flexible solution for CNN based image denoising. TIP, 2018
73. Zhang, Y., Tian, Y., Kong, Y., Zhong, B., Fu, Y.: Residual dense network for image restoration. TPAMI, 2020
74. Zheng, B., Pan, X., Zhang, H., Zhou, X., Slabaugh, G., Yan, C., Yuan, S.: Domainplus: Cross transform domain learning towards efficient high dynamic range imaging. In: ACM MM, 2022
75. Zheng, B., Yuan, S., Slabaugh, G., Leonardis, A.: Image demoreing with learnable bandpass filters. In: CVPR, 2020
76. Zheng, B., Yuan, S., Yan, C., Tian, X., Zhang, J., Sun, Y., Liu, L., Leonardis, A., Slabaugh, G.: Learning frequency domain priors for image demoreing. TPAMI, 2021
77. Zhu, L., Fu, C.W., Lischinski, D., Heng, P.A.: Joint bi-layer optimization for single-image rain streak removal. In: ICCV, 2017
78. Zhu, Q., Mai, J., Shao, L.: A fast single image haze removal algorithm using color attenuation prior. TIP, 2015
79. Zhu, S.C., Mumford, D.: Prior learning and gibbs reaction-diffusion. TPAMI, 1997
80. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. In: ICLR, 2021