

# Navigating Beyond Dropout: An Intriguing Solution Towards Generalizable Image Super Resolution

Hongjun Wang<sup>1\*</sup>, Jiyuan Chen<sup>2\*</sup>, Yinqiang Zheng<sup>1†</sup>, Tiejong Zeng<sup>3†</sup>

<sup>1</sup>The University of Tokyo <sup>2</sup>The Hong Kong Polytechnic University <sup>3</sup>The Chinese University of Hong Kong  
{hjwang-ai@g.ecc, yqzheng@ai}.u-tokyo.ac.jp, jiyuan.chen@connect.polyu.hk, zeng@math.cuhk.edu.hk

<https://github.com/Dreamzz5/Simple-Align>

## Abstract

Deep learning has led to a dramatic leap on Single Image Super-Resolution (SISR) performances in recent years. While most existing work assumes a simple and fixed degradation model (e.g., bicubic downsampling), the research of Blind SR seeks to improve model generalization ability with unknown degradation. Recently, Kong et al. [37] pioneer the investigation of a more suitable training strategy for Blind SR using Dropout [63]. Although such method indeed brings substantial generalization improvements via mitigating overfitting, we argue that Dropout simultaneously introduces undesirable side-effect that compromises model's capacity to faithfully reconstruct fine details. We show both the theoretical and experimental analyses in our paper, and furthermore, we present another easy yet effective training strategy that enhances the generalization ability of the model by simply modulating its first and second-order features statistics. Experimental results have shown that our method could serve as a model-agnostic regularization and outperforms Dropout on seven benchmark datasets including both synthetic and real-world scenarios.

## 1. Introduction

Riding on the waves of the explosive advancement of deep learning technology, Single Image Super-Resolution (SISR) with deep neural networks (DNNs) has greatly evolved in recent years (e.g., VDSR [35], SRResNet [38], EDSR [45], RDN [83] and SwinIR [43]), offering superior performances over traditional prediction models [21, 24, 31, 58].

However, due to the non-trivial issue of collecting a massive amount of natural low-resolution (LR) and high-resolution (HR) image pairs for DNN's training [46, 67], early researchers of SISR resort to manually designed HR/LR image pairs (i.e., bicubic) as a surrogate. Neverthe-




Figure 1. Given a HR image  $I^{HR}$  and its LR version  $I^{LR}$  in part (a), the visual comparisons of the restored results of SRResNet [38], regularized by Dropout [37] and our method respectively, are shown in part (b). Part (c) presents the estimated residuals. **We could observe that our method gives better visual quality and preserves more vivid details.**

less, realistic degradations barely obey such simple assumption, leading to severe performance drop for these models.

Blind SR [32, 46, 65, 73, 80], as an answer to the above question, seeks to improve model generalization ability with unknown degradations. Despite promising results have been obtained with enriching training degradation space (e.g., through handcrafted synthesis [57, 67, 80] or data distribution learning [6, 13, 40]), and enhancing model capability (e.g., unfolding degradation model [30, 86] or exploring image internal statistics [59]), we notice that the in-

\*Equal Contribution

†Corresponding Author

investigation of training strategy (regularization) that benefits Blind SR has been barely touched so far.

In this paper, we argue that such investigation is necessary and meaningful, given current research status of Blind SR. The reasons are as follows. Firstly, without enlarging degradation space for training, the development of Blind SR has reached its bottleneck [40, 46]. Even though there are methods trying to excavate image’s internal similar patterns to perform zero-shot learning [14, 59], they can easily fail in natural cases where self-repeated patterns are absent. On the other hand, training with a large degradation space theoretically grants the model better generalization ability by encouraging it to focus more on learning the shape and texture prior of natural images. Therefore, constructing a diverse degradation pool for training (namely, data-driven based Blind SR, according to Li et al. [40]) has become a mainstream direction for recent Blind SR researches [3, 6, 12, 13, 32, 69, 88], and has been shown effective in both CNN-based [44, 78], GAN-based [68, 80] and Diffusion-based [57, 72] models.

However, only under ideal circumstances that the model trained with diverse degradations can automatically unleash its full potential to learn degradation invariant representation [41], and thus becomes more generalizable to unknown degradations. Experiences from other fields [7, 27, 54, 85] warn that without proper regularization, such ideal case might not be so easily achieved. Recently, Kong et al. [37] first noticed this problem and refreshed the usage of Dropout to mitigate the “overfitting to degradation” issue in data-driven Blind SR. They pointed out that without proper regularization, simply increasing the data and network scale can not continuously improve generalization ability now. Nevertheless, we notice that despite the enhanced performances, Dropout also introduces unwanted side-effects that reduce feature interaction and diversity, which further leads to the loss of vivid high frequency details. We show a preliminary example in Fig. 1, and we will elaborate the theoretical and experimental analyses in Sec. 3.

Moreover, the “overfitting to degradation” problem stems from the network’s excessive attention on some specific degradations. To tackle this issue, we further propose a statistical alignment method that during training, aligns the first and second order feature statistics (i.e., mean and covariance) of two images which have the same content but different degradations. We observe that such simple regularization can effectively enhance the model’s ability to selectively remove degradation-related information during forward pass. Therefore, the model can recover fine-grained HR contents free from distraction of degradations and becomes more generalizable. Our regularization can easily cope with existing popular DNNs in a model-agnostic way and its implementation pre-request (i.e., images with the same content but different degradations) co-

operates smoothly with current data-driven Blind SR methods which have stochastic degradation-generation models [3, 49, 67, 88]. The details of our method will be introduced in Sec. 4 and we here summarize the main contributions of this paper as follows:

- We argue and verify both theoretically and experimentally that Dropout is not a desirable regularization choice for Blind SR setting due to its side-effect in reducing feature interaction and diversity, which further leads to losing information especially in high frequency details.
- We propose a simple statistical alignment method that encourage the model to be thoroughly unaware of degradation information, therefore excavating the full potential of model generalization ability. Note that our regularization actually works in parallel and serves as a complement to existing data-driven Blind SR researches.
- We conduct extensive experiments on seven widely used benchmarks to validate our proposals and arguments.

## 2. Related Work

Blind image super-resolution aims to effectively restore HR images from their LR counterparts with unknown degradations. Over the years, the solutions of this question can be roughly categorized into three brunches. The first brunch of researchers seek to collect real-world HR-LR image pairs for training [8, 11, 82] by adapting the focal length of cameras. However, such collection process is cumbersome and prone to spatial misalignment, making the hope of building a large and diverse training set nearly impossible.

Considering the difficulties faced by the abovementioned work, the second brunch of researchers completely remove the need for external data by performing zero-shot learning. Representative work in this direction includes ZSSR [59] and DGDML-SR [14], which utilize bicubic downsampling and depth information as super-resolution priors respectively. However, these methods rely heavily on the frequently recurring contents of the image, constraining their favourable performances to a very limited set of data.

The core of the third brunch of work lies in enriching the training degradation space, either with hand-crafted synthesis [57, 67, 80] or data distribution learning [6, 13, 40]. Such idea dates back to the very beginning of Blind SR research (i.e., SRMD [78]), and is much in line with the instincts of machine learning that large training space naturally leads to better generalization. Zhang et al. [80] and Wang et al. [67] proposed to use repeated synthesized degradations instead of single ones to build more generic datasets. Later on, to further stimulate real-world degradations, GAN [3, 6, 13, 40, 48, 88] and Diffusion models [72] are incorporated to learn more realistic distributions. For the stochasticity of degradation learning, Bulat et al. [6] and Maeda [50] propose to integrate random vectors into degradation modeling, and Luo et al. [49] further design a unified

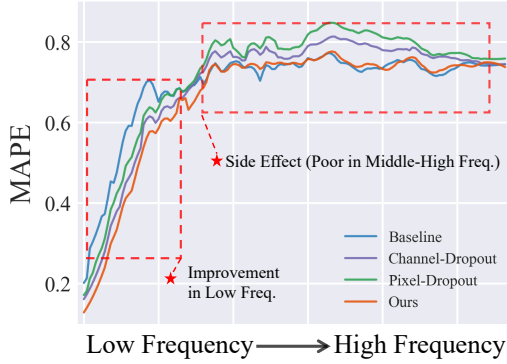


Figure 2. **The MAPE of SRResNet in frequency domain.** We can observe that both channel and pixel dropout have inferior performances to the pure SRResNet in middle-high frequency band, thus losing representation in fine details. On the contrary, our method improves the performances without such side-effect.

probabilistic framework that has general applicability. We also categorize those work that explores model designs on the top of such multi-degradation settings into this brunch, including degradation-adaptive networks [25, 35, 44, 65] and deep unfolding networks [30, 79, 86].

However, as Liu et al. [47] point out, even trained with a large degradation pool, networks still have the tendency to overfit some specific degradations, effectively embedding degradation-related “semantics” within the network and causing reduced generalization. From this perspective, existing works that still adhere to straightforward optimization urgently need a proper training strategy (regularization) that helps to **make the best of the generalization knowledge hidden within the diverse degradations of training data**. Recently, Kong et al. [37] make the first attempt to regularize networks with Dropout [62] and achieve appealing results. However, we argue that Dropout also brings side-effects that hurt high-frequency details in the restored image, leaving rooms for further improvements. Our method circumvents such problem by simply modulating feature statistics to encourage the model becomes indifferent to degradation-aware information, thus squeezing the last bit of degradation-invariant information in training data and improving generalization ability.

With much solid and inspiring work, researches based on the idea of the third brunch has gradually become the mainstream direction of recent Blind SR. Our method in fact serves as a complement to this line of work. In the future even with a larger and more realistic degradation pool, models trained in a straightforward manner still remain susceptible to the possibility of overfitting (e.g., some degradations are easier to learn than others), thus limiting their explorations to degradation-invariant representation (i.e., the ultimate goal of Blind SR). Therefore, our effort actually contributes to the research in a different way, neither from model design nor from building better dataset,

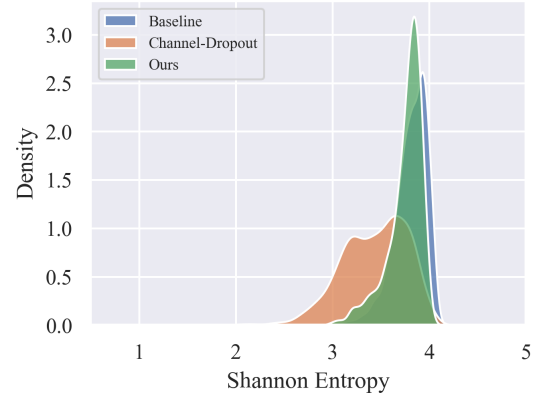


Figure 3. **Comparisons of the channel diversity from frequency perspective.** A higher entropy indicates a wider range of frequency bands covered by the model. Note that since it’s hard to evaluate a pixel from frequency domain, we only consider from channel dimension and investigate the channel-wise dropout here.

but from proposing a training strategy (regularization) that could benefit both existing and future work.

### 3. The Side-effect of Dropout

Dropout [60] and its variants [23, 62] are fundamental techniques used in many high-level vision tasks (e.g., classification) to alleviate co-adaptation and overfitting. However, every coin has its two sides, and the price for Dropout lies in the reduced feature interaction and diversity of the model. While this pose nearly no threat to high-level vision tasks, it can severely impact the performances of image restoration. In this section, we first prove that Dropout decreases feature interactions, and then two experiments are conducted to support our point of view.

Instead of considering the input variables working independently, DNNs encode the interaction between variables for better clues of inferences. For example, the restoration of a human face can be explained as the interactions between left and right cheeks, between eyes and brows, etc. The interaction can be understood as follows. Still with the case of face restoration, let  $\phi_{i=\text{left-cheek}}$  quantifies the numerical importance of the low-resolution left cheek region  $i$  to its high-resolution counterpart when attempting to restore. Then, the interaction utility between the left cheek region  $i$  and the right cheek region  $j$  is measured as the change of  $\phi_{i=\text{left-cheek}}$  value by the presence or absence of the right cheek region  $j$ . If the presence of  $j$  in LR image increases  $\phi_{i=\text{left-cheek}}$  by 0.1 (e.g., the left cheek can interact with right cheek to restore the similar shape and color), we then consider the utility of the interaction between  $i$  and  $j$  to be 0.1.

We further extend the definition of interaction into multi-orders [76] to measure interactions of different complexities. Let  $N$  and  $S$  denote all input units and the context of interaction respectively, and  $T \subseteq S$  is a subset that

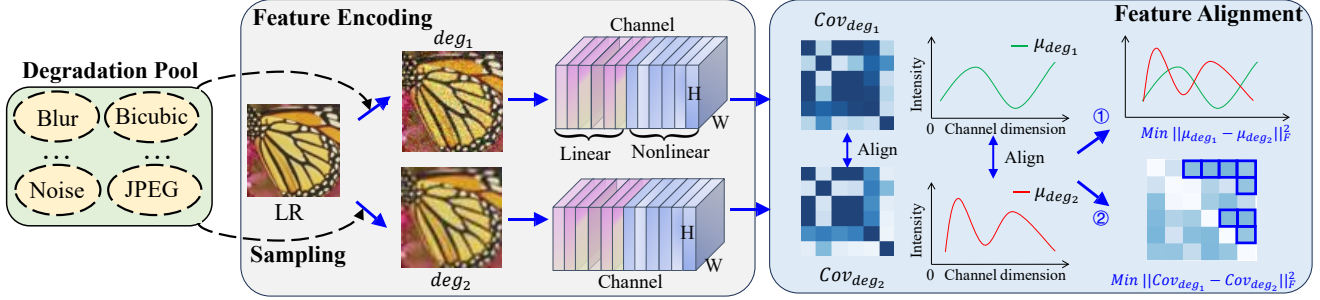


Figure 4. **Overall schema of our proposed regularization.** Given two images with the same content but different degradations (e.g., blur and noise). We compute the mean and covariance of their features respectively (in both linear and nonlinear manner as we will elaborate in Sec. 5). Then, a simple alignment is performed as regularization to encourage the model to learn more natural image prior without the disturbance of degradation information. In this way, the model will be more robust and generalizable against realistic unknown scenarios.

indicates a specific pattern  $T \cup \{i, j\}$ . For any given  $T$ ,  $R^T(i, j)$  quantifies the marginal reward obtained from the pattern of  $T \cup \{i, j\}$ . Therefore we have the  $s$ -order interaction  $I^{(s)}(i, j)$  that measures the average interaction between variables  $(i, j)$  under all possible context with  $s$  units as:

$$I^{(s)}(i, j) = \mathbb{E}_{S \subseteq N \setminus \{i, j\}, |S|=s} \left[ \sum_{T \subseteq S} R^T(i, j) \right]$$

However, when the variables in context  $S$  are randomly removed by Dropout, the computation of  $I_{\text{dropout}}^{(s)}(i, j)$  will only involve the context in  $S' \subseteq S$  that are undropped:

$$I_{\text{dropout}}^{(s)}(i, j) = \mathbb{E}_{S \subseteq N \setminus \{i, j\}, |S|=s} \left[ \mathbb{E}_{S' \subseteq S, |S'|=r} \left( \sum_{T \subseteq S'} R^T(i, j) \right) \right]$$

**Lemma 3.1** *When dropout is applied at a rate  $(1 - p)$ , the interaction  $I_{\text{dropout}}^{(s)}(i, j)$  only comprises of rewards from patterns with at most  $r \sim B(s, p)$  units. Given this, Zhang et al. [77] proved that:  $\frac{I_{\text{dropout}}^{(r)}(i, j)}{I^{(r)}(i, j)} = \frac{\sum_{0 \leq q \leq r} \binom{r}{q} J^{(q)}(i, j)}{\sum_{0 \leq q \leq s} \binom{s}{q} J^{(q)}(i, j)} \leq 1$ , where  $J^q(i, j) = \mathbb{E}_{T \subseteq N \setminus \{i, j\}, |T|=q} [R^T(i, j)]$  is the average interaction for  $(i, j)$  for all potential  $T$  with  $|T| = q$ .*

The above derivation shows that when Dropout is applied to the model, the interaction of every order has become smaller. Extremely speaking, that means the left cheek region will never interact and gain information from the right cheek again, thus undermining its restored quality compared to models with high feature interaction. Note that the proof can also easily extend to channel dimension (i.e., channels work independently without interaction).

There are already many studies emphasizing the importance of interaction, both intra-channel [2, 13, 75, 87] and inter-channel [13, 16–18, 74], in the model’s representation ability of high-frequency components (i.e., fine details such as edges and lines [10, 20]). Therefore, we conjecture that in SISR, applying Dropout will have negative impacts on fine-detail (high frequency) recovery, no matter it is the

pixel-wise Dropout in [60] or the channel-wise one used in [37]. We show some simple observations here, and further provides the LPIPS result which is related to the high-frequency perceptual details [15, 33] in Table 2 of Sec. 5.

We visualize the error of the SRResNet model from the frequency perspective in Fig. 2. The model is trained on DIV2K with different strategies including the one proposed in this paper (see Sec. 4), and tested on the six benchmark datasets with the settings of Real-ESRGAN [66] (i.e., Set5, Set14, BSD100, Test2k, Urban100, and Manga109). The error is estimated by transforming the images into frequency domain via Fast Fourier Transform (FFT). Then we introduce the Mean Absolute Percentage Error (MAPE) [19] metric to indicate the error of each frequency band due to the imbalanced quantity between low and high frequency. Note that here larger MAPE values indicates greater error.

From Fig. 2, we observe that as expected, model trained with Dropout has poorer performances in high frequency recovery. It’s noteworthy that approximately 90% of an image comprises low-frequency signals [70], and human perception is naturally sensitive to high-frequency details of an image. Thus, losing the ability for high frequency restoration usually leads to unsatisfactory perception quality.

Moreover, Dropout also tends to reduce feature diversity by smoothing out the activations of the network like a low-pass filter [84]. Nevertheless, it is acknowledged that the diversity of features actually contributes to the representation ability of different frequency information [51, 55]. In the case of SR, as previously mentioned, low-frequency signals dominate the natural images. Therefore while the network trained with Dropout has no choice but to concentrate the representation power of features into low-frequency (i.e., the features are not diverse enough to represent a wide range of frequency information), network trained without Dropout enjoys diverse feature representation ability in different frequencies. We conduct an auxiliary experiment in Fig. 3, which uses Discrete Cosine Transform as in [55] to first

identify the representative frequency band for each channel, and then computes the Shannon Entropy across channels to reveal the range of frequency band information encoded by the model. The experiment is run with SRResNet on the previously mentioned six datasets, and we present the averaged results here. Without surprise, we observe that model trained with Dropout covers a smaller range of frequency band, limiting its restoration power outside this range.

#### 4. Simple Alignments as Regularization

We revealed the drawbacks of applying Dropout in SR (i.e., reducing feature interactions and diversity) in the previous section. In this section, we will show how a simple alignment can effectively improve the performance of Blind SR. The overall schema of our method is shown in Fig. 4.

The idea of our method stems from the aspiration that model should make predictions independent of different degradations. For example, given two images with the same content but different degradations  $deg_1$  and  $deg_2$ , the model is expected to output the same restored image  $I_o$  from these two inputs, i.e.,  $P(I_o|I_d, d = deg_1) = P(I_o|I_d, d = deg_2)$ . While this seems straightforward by just forcing their intermediate features to be exactly the same, we argue that it will be too harsh and overly constrain the model, hindering its ability to reach a local minimum effectively (*An ablation study showing its inferior performances is in supplemental material*). Instead, in this paper we draw inspirations from image style transfer [29, 42, 64] and treat images with different degradations as with different styles. Note that the similar idea is also adopted in [40], but they are focusing on degradation generation and thus not in our discussion scope.

Then, we follow the traditions of style transfer and utilize the mean and covariance as the degradation (style) sensitive indicators [29]. Although such choice seems lacking solid theoretical foundation, it actually aligns with the research instincts: mean and covariance are two commonly used first and second order statistics in image processing, and have been shown to reflect the global status of activations and the detailed structure and texture respectively [18, 22, 34]. We hypothesize that different degradations should have different impacts on these aspects, therefore making the use of these two statistics reasonable. Empirical studies in Sec. 5 also provide strong evidence for their effectiveness. Consequently, by aligning these statistics across images that have the same content but different degradations, we aim to guide the model to automatically ignore the information specific to degradation during feature encoding, thereby improving its learning ability of degradation-invariant features that are crucially needed for handling new and unknown degradations. Next, we show how exactly the alignment is carried out in both linear and nonlinear manners.

**Linear Alignment** is in its most imaginable form which punishes the statistical differences with Frobenius norm.

Given features  $\mathbf{x}$  and  $\mathbf{x}'$  from images with varying degradation but identical content, the regularization is defined as:

$$\ell_{lin.} = \|Cov(\mathbf{x}) - Cov(\mathbf{x}')\|_F^2 + \|\mu(\mathbf{x}) - \mu(\mathbf{x}')\|_F^2,$$

where  $\|\cdot\|_F^2$  represents the squared matrix Frobenius norm. The mean  $\mu(\mathbf{z}) = \frac{1}{hw} \sum_i \sum_j \mathbf{z}_{:,i,j}$  and covariance  $Cov$  is  $Cov(\mathbf{z}) = \frac{1}{C-1} (\mathbf{z}^\top \mathbf{z} - \frac{1}{C} (\mathbb{1}^\top \mathbf{z})^\top (\mathbb{1}^\top \mathbf{z}))$ .

**Nonlinear Alignment** can be considered as an enhanced version of the linear alignment to overcome the dimensional constraints of encoder models in SR. These models, due to practicality, cannot be arbitrarily scaled up. However theoretically, as feature dimension increases, distributions indistinguishable in low dimensions might become separable in higher ones (i.e., nonlinear property can gradually becomes linearly observable in higher dimensions) [5]. Therefore, we would like to extend the alignment beyond the dimensional limit of the model to align the distributions of different degradations both within the dimensional capacity of the model and beyond it in higher or even infinite dimension. We believe the knowledge glimpsed from the higher dimension can effectively “highlight” the deeper differences between degradations and thus, serves as a “look-ahead” to further steer the model towards degradation-invariant. Meanwhile, the gradients from higher-dimensional alignment might also act like a perceptive teacher to facilitate the optimization of alignment in lower dimensions.

To this end, in order to peep secrets from the higher dimension in a parameter-efficient way, let’s consider a feature mapping  $\Phi: \mathcal{X} \rightarrow \mathcal{H}$ , which transforms elements in  $\mathcal{X}$  into a reproducing kernel Hilbert space (RKHS) denoted as  $\mathcal{H}$ . Instead of deriving the explicit expression of  $\Phi$ , in this paper we utilize the Random Fourier features (RFF) [56] to approximate the behavior of Radial Basis Function (RBF) kernel to explicitly project features into  $\mathcal{H}_{RFF}$  where data has limited dimension but shows similar properties as if in infinite dimension space. Therefore, let  $\mathbf{u}$  and  $\mathbf{v}$  be the RFF mapping function satisfying  $\mathbf{u}(z), \mathbf{v}(z) \in \mathcal{H}_{RFF}$ , we have our nonlinear alignment as:

$$\begin{aligned} \ell_{nonlin.} &= \|Cov(\mathbf{h}_x) - Cov(\mathbf{h}_{x'})\|_F^2 + \|\mu(\mathbf{h}_x) - \mu(\mathbf{v}(\mathbf{h}_{x'}))\|_F^2, \\ \mathbf{h}_x &= (\mathbf{u}^1(x), \dots, \mathbf{u}^n(x)), \quad \mathbf{h}_{x'} = (\mathbf{v}^1(x'), \dots, \mathbf{v}^n(x')), \\ \mathcal{H}_{RFF} &= \left\{ h: \sqrt{2} \cos(\omega x + \phi) \mid \omega \sim \mathcal{N}(0, 1), \phi \sim U(0, 2\pi) \right\}. \end{aligned}$$

**Discussion.** Thanks to the advancement of realistic stochastic degradation generation methods, training models with multi-degradations has become standard for most recent Blind SR works. Our method integrates seamlessly with them by simply requiring the degradation generation model (e.g., [49, 68]) to randomly generate one more degradation for each (or some) image(s) to form the  $(\mathbf{x}, \mathbf{x}')$  pairs for regularization. Note that the forward pass of  $(\mathbf{x}, \mathbf{x}')$  can be efficiently parallelized across multiple GPUs. Furthermore,

Table 1. Six benchmarks with eight types of degradations (clean, noise, blur, jpeg, blur+noise, blur+jpeg, noise+jpeg, and blur+noise+jpeg) are used to evaluate the PSNR (dB) results in  $\times 4$  resolution settings. Degradations are abbreviated in table.

Models	Set5 [4]				Set14 [71]				BSD100 [52]			
	clean	blur	noise	jpeg	clean	blur	noise	jpeg	clean	blur	noise	jpeg
SRResNet [38]	24.85	24.73	22.52	23.67	23.25	23.05	21.18	22.32	23.06	22.99	21.34	22.47
+ Ours	25.93	25.62	23.15	24.38	24.12	23.80	21.67	22.99	23.83	23.64	21.77	23.04
Improvement	+1.08	+0.89	+0.63	+0.71	+0.87	+0.75	+0.49	+0.67	+0.77	+0.65	+0.43	+0.57
RRDB [66]	25.18	25.12	21.79	23.82	23.74	23.36	21.02	22.59	23.38	23.32	21.00	22.73
+ Ours	26.78	26.55	23.02	24.70	24.70	24.35	21.91	23.21	24.59	24.54	23.47	23.67
Improvement	+1.60	+1.43	+1.23	+0.88	+0.96	+0.99	+0.89	+0.62	+1.21	+1.22	+2.47	+0.94
MSRN [39]	25.25	24.89	22.57	24.08	23.38	23.10	21.80	22.53	23.38	23.30	21.92	22.76
+ Ours	25.81	25.52	22.84	24.46	23.93	23.64	21.86	22.83	23.72	23.58	22.01	22.98
Improvement	+0.56	+0.63	+0.27	+0.38	+0.55	+0.54	+0.06	+0.30	+0.34	+0.28	+0.09	+0.22
SwinIR [43]	26.25	26.03	22.96	24.37	24.53	24.25	22.08	23.14	23.91	23.83	22.12	23.04
+ Ours	26.49	26.23	24.61	24.68	24.65	24.28	22.23	23.29	24.04	23.96	22.21	23.15
Improvement	+0.24	+0.20	+1.65	+0.31	+0.12	+0.03	+0.15	+0.15	+0.13	+0.13	+0.09	+0.11

	b+n	b+j	n+j	b+n+j	b+n	b+j	n+j	b+n+j	b+n	b+j	n+j	b+n+j
SRResNet [38]	23.27	23.40	23.05	22.73	22.23	22.06	21.99	21.77	22.25	22.33	22.22	22.04
+ Ours	23.79	23.86	23.71	23.19	22.65	22.63	22.55	22.16	22.53	22.79	22.62	22.32
Improvement	+0.52	+0.46	+0.66	+0.46	+0.42	+0.57	+0.56	+0.39	+0.28	+0.46	+0.40	+0.28
RRDB [66]	23.44	23.45	23.32	22.81	22.47	22.17	22.29	21.95	22.39	22.47	22.42	22.15
+ Ours	24.12	24.14	23.93	23.26	22.80	22.76	22.71	22.21	22.85	23.21	22.97	22.54
Improvement	+0.68	+0.69	+0.61	+0.45	+0.33	+0.59	+0.42	+0.26	+0.46	+0.74	+0.55	+0.39
MSRN [39]	23.55	23.59	23.50	22.95	22.39	22.23	22.19	21.97	22.57	22.61	22.45	22.24
+ Ours	23.70	23.80	23.73	23.06	22.52	22.49	22.48	22.08	22.68	22.73	22.56	22.26
Improvement	+0.15	+0.21	+0.23	+0.11	+0.13	+0.26	+0.29	+0.11	+0.11	+0.12	+0.11	+0.02
SwinIR [43]	23.80	23.84	23.67	22.99	22.53	22.73	22.59	22.20	22.61	22.82	22.61	22.34
+ Ours	24.13	24.17	23.89	23.09	22.87	22.79	22.81	22.28	22.77	22.98	22.76	22.40
Improvement	+0.33	+0.33	+0.22	+0.10	+0.34	+0.06	+0.22	+0.08	+0.16	+0.16	+0.15	+0.06

Models	Test2k [36]				Urban100 [28]				Manga109 [53]			
	clean	blur	noise	jpeg	clean	blur	noise	jpeg	clean	blur	noise	jpeg
SRResNet [38]	23.91	23.71	21.77	23.11	21.23	21.06	19.74	20.60	18.42	18.75	18.08	18.27
+ Ours	24.58	24.43	22.17	23.65	21.94	21.65	20.19	21.20	19.18	19.46	18.90	19.02
Improvement	+0.67	+0.72	+0.40	+0.54	+0.71	+0.59	+0.45	+0.60	+0.76	+0.71	+0.82	+0.75
RRDB [66]	24.16	23.64	21.34	23.36	21.57	21.18	19.61	20.93	18.59	18.64	18.30	18.41
+ Ours	24.97	24.76	22.15	23.86	22.29	21.95	20.21	21.40	19.40	19.61	18.96	19.24
Improvement	+0.81	+1.12	+0.81	+0.50	+0.72	+0.77	+0.60	+0.47	+0.81	+0.97	+0.66	+0.83
MSRN [39]	22.99	23.83	22.30	23.22	21.35	21.14	20.19	20.75	19.12	19.31	18.72	18.89
+ Ours	24.52	24.23	22.38	23.56	21.88	21.54	20.22	21.14	19.23	19.35	18.84	19.01
Improvement	+1.53	+0.40	+0.08	+0.34	+0.53	+0.40	+0.03	+0.39	+0.11	+0.04	+0.12	+0.12
SwinIR [43]	24.78	24.57	22.71	23.63	22.18	21.90	20.56	21.32	19.10	19.27	18.71	18.95
+ Ours	24.98	24.76	22.84	23.80	22.34	22.07	20.69	21.48	19.24	19.45	18.98	19.28
Improvement	+0.20	+0.19	+0.13	+0.17	+0.16	+0.17	+0.13	+0.16	+0.14	+0.18	+0.27	+0.33

	b+n	b+j	n+j	b+n+j	b+n	b+j	n+j	b+n+j	b+n	b+j	n+j	b+n+j
SRResNet [38]	22.81	22.87	22.85	22.59	20.46	20.30	20.42	20.09	18.59	18.50	18.21	18.39
+ Ours	23.11	23.27	23.22	22.82	20.73	20.72	20.91	20.37	19.27	19.17	18.98	19.01
Improvement	+0.30	+0.40	+0.37	+0.23	+0.27	+0.42	+0.49	+0.28	+0.68	+0.67	+0.77	+0.62
RRDB [66]	22.93	22.87	23.12	22.73	20.57	20.40	20.74	20.24	18.83	18.43	18.38	18.41
+ Ours	23.14	23.37	23.34	22.82	20.76	20.85	21.03	20.38	19.43	19.31	19.12	19.15
Improvement	+0.21	+0.50	+0.22	+0.09	+0.19	+0.45	+0.29	+0.14	+0.60	+0.88	+0.74	+0.74
MSRN [39]	23.03	23.01	22.94	22.66	20.65	20.43	20.56	20.19	19.16	19.02	18.80	18.88
+ Ours	23.21	23.21	23.18	22.78	20.76	20.64	20.89	20.26	19.19	19.18	18.92	18.93
Improvement	+0.18	+0.20	+0.24	+0.12	+0.11	+0.21	+0.33	+0.07	+0.03	+0.16	+0.12	+0.05
SwinIR [43]	23.15	23.27	23.21	22.81	20.89	20.79	20.98	20.45	19.07	19.02	18.79	18.80
+ Ours	23.35	23.47	23.45	22.93	21.02	20.98	21.12	20.53	19.37	19.35	19.15	19.12
Improvement	+0.20	+0.20	+0.24	+0.12	+0.13	+0.19	+0.14	+0.08	+0.30	+0.33	+0.36	+0.32




Figure 5. Visual results of representative degradations (blur, noise, and JPEG) on baseline models w/ and w/o our regularization.

explicitly aligning only the feature’s first and second order statistics also makes our method reliable and efficient.

## 5. Experiment

In this paper we follow Kong et al. [37] and adopt the widely acknowledged multi-degradations settings used in Blind SR researches [68] to make fair and credible comparisons. *Due to space limits, we refer readers to our supplementary materials for more experimental settings and ablation studies.*

**Improvements on Baseline Models.** We show the improvements of applying our methods on the top of baseline models in Table 2. Notably, we achieve considerable performance improvements in almost all cases (averaged 0.44dB). We draw three conclusions from Table 2: (1) the test case of blur + noise + jpeg (b+n+j) has universally the smallest improvement (averaged 0.21dB) to baseline models compared with other degradations types (averaged 0.45dB). We hypothesize that this is because b+n+j is actually an “in-

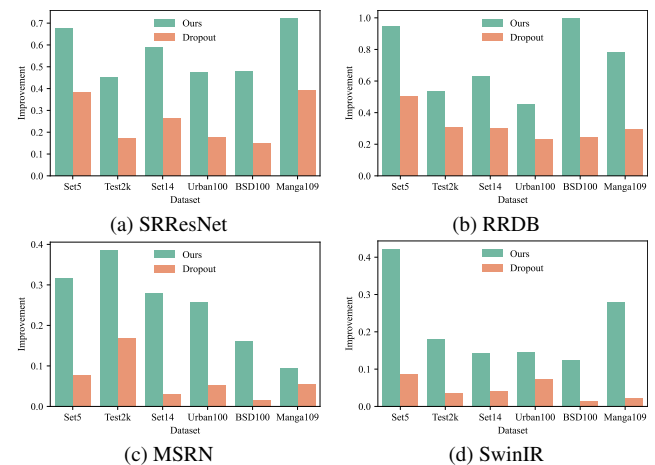


Figure 6. Average improvements on baseline models of our method and Dropout on six benchmark datasets.

Table 2. Comparisons with Dropout using both PSNR (run on realistic data of NTIRE 2018 challenge) and LPIPS (run on six widely used benchmarks).

Models	PSNR $\uparrow$			LPIPS $\downarrow$					
	mild / difficult / wild	Set5 / Set14 / BSD / Manga / Urban / Test2k							
SRResNet	16.94 / 17.84 / 17.55	0.241 / 0.353 / 0.433 / 0.220 / 0.323 / 0.369							
+Dropout	17.10 / 18.02 / 17.78	0.245 / 0.356 / 0.439 / 0.223 / 0.333 / 0.374							
+Ours	<b>17.69 / 18.30 / 17.94</b>	<b>0.238 / 0.349 / 0.429 / 0.219 / 0.319 / 0.367</b>							
RRDB	16.79 / 17.66 / 17.38	0.218 / 0.324 / 0.398 / 0.206 / 0.300 / 0.344							
+Dropout	17.31 / 18.17 / 17.89	0.229 / 0.334 / 0.412 / 0.210 / 0.309 / 0.354							
+Ours	<b>17.88 / 18.47 / 18.21</b>	<b>0.217 / 0.322 / 0.397 / 0.205 / 0.299 / 0.341</b>							




Figure 7. Restoration results of Dropout and our method.

distribution” test setting where the models already have good enough performances. We call it “in-distribution” because following Wang et al. [68], the model is also (and only) trained with such setting, although with different degradation parameters. Since our method is designed to improve performances over unknown degradations, it is not surprising that we don’t have impressive results with “in-distribution” setting. (2) On the contrary, the test cases of clean, blur, noise and jpeg, whose patterns deviate the most from training distribution (i.e.,  $b+n+j$ ), have shown excellent improvements (averaged 0.55dB) compared with other settings (averaged 0.31dB), demonstrating the improved generalization ability offered by our method. (3) From the perspective of models, SRResNet and RRDB show more improvements (averaged 0.65dB) than MSRN and SwinIR (averaged 0.22dB). We speculate that it’s because MSRN and SwinIR introduce extra multi-scale and transformer architecture over the residual CNNs of SRResNet and RRDB, suggesting that multi-scale feature and larger receptive field might also benefit model generalization in Blind SR. We also show the visual comparisons of baseline models w/ and w/o our regularization in Fig. 5, and as observed, models regularized by our method generally perform better in content reconstruction and artifact removal.

**Performance Comparison with Dropout.** The compari-

son of average improvement between Dropout (unless otherwise specified, Dropout used in our experiments refers to the channel-wise Dropout used in [37]) and our method is illustrated in Fig. 6. It is noteworthy that our method outperforms Dropout on all models and datasets, with an average improvement of 0.57dB v.s. 0.26dB on SRResNet, 0.72dB v.s. 0.31dB on RRDB, 0.25dB v.s. 0.07dB on MSRN and 0.22dB v.s. 0.05dB on SwinIR. We argue that the performance gap observed here is reasonable, because our method in essence explicitly aligns different degradations to encourage the model becomes degradation-invariant. However, Dropout, aside from its side-effects discussed earlier, only guides such invariant feature learning in a rather implicit manner, and thus can’t thoroughly get rid of degradation-specific information, leading to inferior generalization ability. *The detailed data of Fig. 6 is in supplementary material.*

Although the multi-degradation setting of [68] has been proven effective in stimulating real-world degradations, we further evaluate our approach against Dropout with realistic NTIRE 2018 SR challenge data [61]. The results are presented in the left part of Table 2 with our method showing overwhelming advantages in generalization over Dropout. Moreover, we also investigate the model performances using the LPIPS metric [81] which is more relevant to perceptual details of an image in the right part of Table 2. The results of LPIPS correspond to our previous discussion that while Dropout improves PSNR by alleviating overfitting, it could simultaneously bring negative impacts on the recovery of fine-grained details (reflecting by LPIPS). On the contrary, our method can enjoy the best of both worlds without such side-effect. We also visualize the restored images of Dropout and our method for vivid comparison in Fig. 7.

## 6. Conclusion

From theoretical and experimental aspects, this paper first reveal the side-effects of applying Dropout as a regularizer in SR. Then as an alternative, we propose a simple and effective feature alignment regularization that can further enhance the generalization ability for Blind SR models. *Given the current challenges in advancing Blind SR researches, we call for more efforts exploring training regularization, a path not fully developed but potentially highly impactful.*

## 7. Acknowledgments

This work was supported by Grant NSFC/RGC N\_CUHK 415/19, Grant ITF ITS/173/22FP, Grant RGC 14300219, 14302920, 14301121, CUHK Direct Grant for Research, JSPS KAKENHI Grant Numbers 22H00529, 20H05951, JST-Mirai Program JPMJMI23G1 and JST SPRING, Grant Number JPMJSP2108.



## References

- [1] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2017. 1
- [2] Jiawang Bai, Li Yuan, Shu-Tao Xia, Shuicheng Yan, Zhifeng Li, and Wei Liu. Improving vision transformers by revisiting high-frequency components. In *European Conference on Computer Vision*, pages 1–18. Springer, 2022. 4
- [3] Sefi Bell-Kligler, Assaf Shocher, and Michal Irani. Blind super-resolution kernel estimation using an internal-gan. *Advances in Neural Information Processing Systems*, 32, 2019. 2
- [4] Marco Bevilacqua, Aline Roumy, Christine Guillemot, and Marie Line Alberi-Morel. Low-complexity single-image super-resolution based on nonnegative neighbor embedding. 2012. 6, 1, 3
- [5] Himani Bhavsar and Mahesh H Panchal. A review on support vector machine for data classification. *International Journal of Advanced Research in Computer Engineering & Technology (IJARCET)*, 1(10):185–189, 2012. 5
- [6] Adrian Bulat, Jing Yang, and Georgios Tzimiropoulos. To learn image super-resolution, use a gan to learn how to do image degradation first. In *Proceedings of the European conference on computer vision (ECCV)*, pages 185–200, 2018. 1, 2
- [7] Remi Cadene, Corentin Dancette, Matthieu Cord, Devi Parikh, et al. Rubi: Reducing unimodal biases for visual question answering. *Advances in neural information processing systems*, 32, 2019. 2
- [8] Jianrui Cai, Hui Zeng, Hongwei Yong, Zisheng Cao, and Lei Zhang. Toward real-world single image super-resolution: A new benchmark and a new model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3086–3095, 2019. 2
- [9] Tadeusz Caliński and Jerzy Harabasz. A dendrite method for cluster analysis. *Communications in Statistics-theory and Methods*, 3(1):1–27, 1974. 2
- [10] Fergus W Campbell and John G Robson. Application of fourier analysis to the visibility of gratings. *The Journal of physiology*, 197(3):551, 1968. 4
- [11] Chang Chen, Zhiwei Xiong, Xinmei Tian, Zheng-Jun Zha, and Feng Wu. Camera lens super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1652–1660, 2019. 2
- [12] Du Chen, Jie Liang, Xindong Zhang, Ming Liu, Hui Zeng, and Lei Zhang. Human guided ground-truth generation for realistic image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14082–14091, 2023. 2
- [13] Xuhai Chen, Jiangning Zhang, Chao Xu, Yabiao Wang, Chengjie Wang, and Yong Liu. Better” cmos” produces clearer images: Learning space-variant blur estimation for blind image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1651–1661, 2023. 1, 2, 4
- [14] Xi Cheng, Zhenyong Fu, and Jian Yang. Zero-shot image super-resolution with depth guided internal degradation learning. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVII 16*, pages 265–280. Springer, 2020. 2
- [15] Vishal Chudasama and Kishor Upla. RsrGAN: computationally efficient real-world single image super-resolution using generative adversarial network. *Machine Vision and Applications*, 32(1):3, 2021. 4
- [16] Kai Cui, Atanas Boev, Elena Alshina, and Eckehard Steinbach. Color image restoration exploiting inter-channel correlation with a 3-stage CNN. *IEEE Journal of Selected Topics in Signal Processing*, 15(2):174–189, 2020. 4
- [17] Yuning Cui and Alois Knoll. Exploring the potential of channel interactions for image restoration. *Knowledge-Based Systems*, 282:111156, 2023.
- [18] Tao Dai, Jianrui Cai, Yongbing Zhang, Shu-Tao Xia, and Lei Zhang. Second-order attention network for single image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 11065–11074, 2019. 4, 5
- [19] Arnaud De Myttenaere, Boris Golden, Bénédicte Le Grand, and Fabrice Rossi. Mean absolute percentage error for regression models. *Neurocomputing*, 192:38–48, 2016. 4
- [20] Russell L De Valois and Karen K De Valois. Spatial vision. *Annual review of psychology*, 31(1):309–341, 1980. 4
- [21] Raanan Fattal. Image upsampling via imposed edge statistics. In *ACM SIGGRAPH 2007 papers*, pages 95–es. 2007. 1
- [22] Shang-Hua Gao, Qi Han, Duo Li, Ming-Ming Cheng, and Pai Peng. Representative batch normalization with feature calibration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8669–8679, 2021. 5
- [23] Golnaz Ghiasi, Tsung-Yi Lin, and Quoc V Le. Dropblock: A regularization method for convolutional networks. *arXiv preprint arXiv:1810.12890*, 2018. 3
- [24] Daniel Glasner, Shai Bagon, and Michal Irani. Super-resolution from a single image. In *2009 IEEE 12th international conference on computer vision*, pages 349–356. IEEE, 2009. 1
- [25] Jinjin Gu, Hannan Lu, Wangmeng Zuo, and Chao Dong. Blind super-resolution with iterative kernel correction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1604–1613, 2019. 3
- [26] Jinjin Gu, Haoming Cai, Haoyu Chen, Xiaoxing Ye, Jimmy Ren, and Chao Dong. Pipal: a large-scale image quality assessment dataset for perceptual image restoration. In *European Conference on Computer Vision*, pages 633–651. Springer, 2020. 1
- [27] Xinqian Gu, Hong Chang, Bingpeng Ma, Shutao Bai, Shiguang Shan, and Xilin Chen. Clothes-changing person re-identification with rgb modality only. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1060–1069, 2022. 2
- [28] Jia-Bin Huang, Abhishek Singh, and Narendra Ahuja. Single image super-resolution from transformed self-exemplars. In

- Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5197–5206, 2015. 6, 1, 3
- [29] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE international conference on computer vision*, pages 1501–1510, 2017. 5
- [30] Yan Huang, Shang Li, Liang Wang, Tieniu Tan, et al. Unfolding the alternating optimization for blind super resolution. *Advances in Neural Information Processing Systems*, 33:5632–5643, 2020. 1, 3
- [31] Michal Irani and Shmuel Peleg. Improving resolution by image registration. *CVGIP: Graphical models and image processing*, 53(3):231–239, 1991. 1
- [32] Xiaozhong Ji, Yun Cao, Ying Tai, Chengjie Wang, Jilin Li, and Feiyue Huang. Real-world super-resolution via kernel estimation and noise injection. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 466–467, 2020. 1, 2
- [33] Zexin Ji, Xin Dong, Zhendong Li, Zekuan Yu, and Hao Liu. Non-local network routing for perceptual image super-resolution. In *Pattern Recognition and Computer Vision: 4th Chinese Conference, PRCV 2021, Beijing, China, October 29–November 1, 2021, Proceedings, Part III 4*, pages 164–176. Springer, 2021. 4
- [34] Levent Karacan, Erkut Erdem, and Aykut Erdem. Structure-preserving image smoothing via region covariances. *ACM Transactions on Graphics (TOG)*, 32(6):1–11, 2013. 5
- [35] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1646–1654, 2016. 1, 3
- [36] Xiangtao Kong, Hengyuan Zhao, Yu Qiao, and Chao Dong. Classsr: A general framework to accelerate super-resolution networks by data characteristic. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12016–12025, 2021. 6, 1, 3
- [37] Xiangtao Kong, Xina Liu, Jinjin Gu, Yu Qiao, and Chao Dong. Reflash dropout in image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6002–6012, 2022. 1, 2, 3, 4, 7, 8
- [38] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017. 1, 6, 3
- [39] Juncheng Li, Faming Fang, Kangfu Mei, and Guixu Zhang. Multi-scale residual network for image super-resolution. In *Proceedings of the European conference on computer vision (ECCV)*, pages 517–532, 2018. 6, 3
- [40] Xiaoming Li, Chaofeng Chen, Xianhui Lin, Wangmeng Zuo, and Lei Zhang. From face to natural image: Learning real degradation for blind image super-resolution. In *European Conference on Computer Vision*, pages 376–392. Springer, 2022. 1, 2, 5
- [41] Xin Li, Bingchen Li, Xin Jin, Cuiling Lan, and Zhibo Chen. Learning distortion invariant representation for image restoration from a causality perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1714–1724, 2023. 2
- [42] Yanghao Li, Naiyan Wang, Jiaying Liu, and Xiaodi Hou. Demystifying neural style transfer. *arXiv preprint arXiv:1701.01036*, 2017. 5
- [43] Jingyun Liang, Jiezhong Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. Swinir: Image restoration using swin transformer. In *IEEE International Conference on Computer Vision Workshops*, 2021. 1, 6, 3
- [44] Jie Liang, Hui Zeng, and Lei Zhang. Efficient and degradation-adaptive network for real-world image super-resolution. In *European Conference on Computer Vision*, pages 574–591. Springer, 2022. 2, 3
- [45] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2017. 1
- [46] Anran Liu, Yihao Liu, Jinjin Gu, Yu Qiao, and Chao Dong. Blind image super-resolution: A survey and beyond. *IEEE transactions on pattern analysis and machine intelligence*, 45(5):5461–5480, 2022. 1, 2
- [47] Yihao Liu, Anran Liu, Jinjin Gu, Zhipeng Zhang, Wenhao Wu, Yu Qiao, and Chao Dong. Discovering distinctive” semantics” in super-resolution networks. *arXiv preprint arXiv:2108.00406*, 2021. 3, 1, 2
- [48] Andreas Lugmayr, Martin Danelljan, and Radu Timofte. Un-supervised learning for real-world super-resolution. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 3408–3416. IEEE, 2019. 2
- [49] Zhengxiong Luo, Yan Huang, Shang Li, Liang Wang, and Tieniu Tan. Learning the degradation distribution for blind image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6063–6072, 2022. 2, 5
- [50] Shunta Maeda. Unpaired image super-resolution using pseudo-supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 291–300, 2020. 2
- [51] Salma Abdel Magid, Yulun Zhang, Donglai Wei, Won-Dong Jang, Zudi Lin, Yun Fu, and Hanspeter Pfister. Dynamic high-pass filtering and multi-spectral attention for image super-resolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4288–4297, 2021. 4
- [52] David Martin, Charless Fowlkes, Doron Tal, and Jitendra Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, pages 416–423. IEEE, 2001. 6, 1, 3
- [53] Yusuke Matsui, Kota Ito, Yuji Aramaki, Azuma Fujimoto, Toru Ogawa, Toshihiko Yamasaki, and Kiyoharu Aizawa. Sketch-based manga retrieval using manga109 dataset. *Mul-*

- timedia Tools and Applications*, 76(20):21811–21838, 2017. 6, 1, 3
- [54] Wei Qin, Hanwang Zhang, Richang Hong, Ee-Peng Lim, and Qianru Sun. Causal interventional training for image recognition. *IEEE Transactions on Multimedia*, 2021. 2
- [55] Zequn Qin, Pengyi Zhang, Fei Wu, and Xi Li. Fcanet: Frequency channel attention networks. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 783–792, 2021. 4
- [56] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. *Advances in neural information processing systems*, 20, 2007. 5
- [57] Hshmat Sahak, Daniel Watson, Chitwan Saharia, and David Fleet. Denoising diffusion probabilistic models for robust image super-resolution in the wild. *arXiv preprint arXiv:2302.07864*, 2023. 1, 2
- [58] Qi Shan, Zhaorong Li, Jiaya Jia, and Chi-Keung Tang. Fast image/video upsampling. *ACM Transactions on Graphics (TOG)*, 27(5):1–7, 2008. 1
- [59] Assaf Shocher, Nadav Cohen, and Michal Irani. “zero-shot” super-resolution using deep internal learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3118–3126, 2018. 1, 2
- [60] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014. 3, 4
- [61] Radu Timofte, Eirikur Agustsson, Luc Van Gool, Ming-Hsuan Yang, and Lei Zhang. Ntire 2017 challenge on single image super-resolution: Methods and results. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 114–125, 2017. 8, 1
- [62] Jonathan Tompson, Ross Goroshin, Arjun Jain, Yann LeCun, and Christoph Bregler. Efficient object localization using convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 648–656, 2015. 3
- [63] Jonathan Tompson, Ross Goroshin, Arjun Jain, Yann LeCun, and Christoph Bregler. Efficient object localization using convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 648–656, 2015. 1
- [64] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6924–6932, 2017. 5
- [65] Longguang Wang, Yingqian Wang, Xiaoyu Dong, Qingyu Xu, Juegang Yang, Wei An, and Yulan Guo. Unsupervised degradation representation learning for blind super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10581–10590, 2021. 1, 3
- [66] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 0–0, 2018. 4, 6, 3
- [67] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1905–1914, 2021. 1, 2
- [68] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *International Conference on Computer Vision Workshops (ICCVW)*, 2021. 2, 5, 7, 8, 1
- [69] Yunxuan Wei, Shuhang Gu, Yawei Li, Radu Timofte, Longcun Jin, and Hengjie Song. Unsupervised real-world image super resolution via domain-distance aware training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13385–13394, 2021. 2
- [70] Tianshuo Xu, Peng Mi, Xiawu Zheng, Lijiang Li, Fei Chao, Guannan Jiang, Wei Zhang, Yiyi Zhou, and Rongrong Ji. What hinders perceptual quality of psnr-oriented methods? *arXiv preprint arXiv:2201.01034*, 2022. 4
- [71] Jianchao Yang, John Wright, Thomas S Huang, and Yi Ma. Image super-resolution via sparse representation. *IEEE transactions on image processing*, 19(11):2861–2873, 2010. 6, 1, 3
- [72] Tao Yang, Peiran Ren, Lei Zhang, et al. Synthesizing realistic image restoration training pairs: A diffusion approach. *arXiv preprint arXiv:2303.06994*, 2023. 2
- [73] Yuan Yuan, Siyuan Liu, Jiawei Zhang, Yongbing Zhang, Chao Dong, and Liang Lin. Unsupervised image super-resolution using cycle-in-cycle generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 701–710, 2018. 1
- [74] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5728–5739, 2022. 4
- [75] Guiwei Zhang, Yongfei Zhang, Tianyu Zhang, Bo Li, and Shiliang Pu. Pha: Patch-wise high-frequency augmentation for transformer-based person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14133–14142, 2023. 4
- [76] Hao Zhang, Xu Cheng, Yiting Chen, and Quanshi Zhang. Game-theoretic interactions of different orders. *arXiv preprint arXiv:2010.14978*, 2020. 3
- [77] Hao Zhang, Sen Li, Yinchao Ma, Mingjie Li, Yichen Xie, and Quanshi Zhang. Interpreting and boosting dropout from a game-theoretic view. *arXiv preprint arXiv:2009.11729*, 2020. 4
- [78] Kai Zhang, Wangmeng Zuo, and Lei Zhang. Learning a single convolutional super-resolution network for multiple degradations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3262–3271, 2018. 2
- [79] Kai Zhang, Luc Van Gool, and Radu Timofte. Deep unfolding network for image super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3217–3226, 2020. 3

- [80] Kai Zhang, Jingyun Liang, Luc Van Gool, and Radu Timofte. Designing a practical degradation model for deep blind image super-resolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4791–4800, 2021. [1](#), [2](#)
- [81] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. [8](#)
- [82] Xuaner Zhang, Qifeng Chen, Ren Ng, and Vladlen Koltun. Zoom to learn, learn to zoom. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3762–3770, 2019. [2](#)
- [83] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2472–2481, 2018. [1](#)
- [84] Zhilu Zhang, Adrian V Dalca, and Mert R Sabuncu. Confidence calibration for convolutional neural networks using structured dropout. *arXiv preprint arXiv:1906.09551*, 2019. [4](#)
- [85] Tong Zhao, Gang Liu, Daheng Wang, Wenhao Yu, and Meng Jiang. Learning from counterfactual links for link prediction. In *International Conference on Machine Learning*, pages 26911–26926. PMLR, 2022. [2](#)
- [86] Hongyi Zheng, Hongwei Yong, and Lei Zhang. Unfolded deep kernel estimation for blind image super-resolution. In *European Conference on Computer Vision*, pages 502–518. Springer, 2022. [1](#), [3](#)
- [87] Man Zhou, Jie Huang, Chun-Le Guo, and Chongyi Li. Fourmer: an efficient global modeling paradigm for image restoration. In *International Conference on Machine Learning*, pages 42589–42601. PMLR, 2023. [4](#)
- [88] Ruofan Zhou and Sabine Susstrunk. Kernel modeling super-resolution on real low-resolution images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2433–2443, 2019. [2](#)