

This CVPR 2021 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Debiased Subjective Assessment of Real-World Image Enhancement

Peibei Cao¹, Zhangyang Wang², and Kede Ma¹ ¹ City University of Hong Kong, ² University of Texas at Austin

peibeicao2-c@my.cityu.edu.hk, atlaswang@utexas.edu, kede.ma@cityu.edu.hk

Abstract

In real-world image enhancement, it is often challenging (if not impossible) to acquire ground-truth data, preventing the adoption of distance metrics for objective quality assessment. As a result, one often resorts to subjective quality assessment, the most straightforward and reliable means of evaluating image enhancement. Conventional subjective testing requires manually pre-selecting a small set of visual examples, which may suffer from three sources of biases: 1) sampling bias due to the extremely sparse distribution of the selected samples in the image space; 2) algorithmic bias due to potential overfitting the selected samples; 3) subjective bias due to further potential cherry-picking test results. This eventually makes the field of real-world image enhancement more of an art than a science. Here we take steps towards debiasing conventional subjective assessment by automatically sampling a set of adaptive and diverse images for subsequent testing. This is achieved by casting sample selection into a joint maximization of the discrepancy between the enhancers and the diversity among the selected input images. Careful visual inspection on the resulting enhanced images provides a debiased ranking of the enhancement algorithms. We demonstrate our subjective assessment method using three popular and practically demanding image enhancement tasks: dehazing, super-resolution, and low-light enhancement.

1. Introduction

For many years, image enhancement has been investigated in an unrealistic setting, with the assumption that the original images of perfect quality exist to help evaluate visual quality of the enhanced images. This promotes the use of full-reference image quality metrics [69] to compute an average distance between a large set of enhanced and original image pairs as an indication of enhancement performance. Along this path, many full-reference metrics have been proposed [71, 54, 82, 9], trying to measure this distance more perceptually.

However, in real-world image enhancement, it is of-

ten difficult (if not impossible) to specify desired outputs. Moreover, there may be multiple diverse outputs that are desirable, as in the case of super-resolution [75]. Therefore, full-reference models that rely on a single "ideal" image are not applicable. Some attempts have been made to adopt noreference models [70] for performance assessment of realworld enhancement. However, no-reference objective assessment is still in its infancy, and accurate models for (specific or general) image enhancement applications are largely lacking. Currently, the most widely used no-reference metric - NIQE [45] - was empirically proven to correlate poorly with human quality judgments of the enhanced images [45], which exhibit unique and algorithm-specific artifacts that are often non-overlapping with natural distortions.

Alternatively, one may refer to subjective quality assessment, which is so far the most straightforward and reliable way of evaluating real-world image enhancement because the ultimate receiver in most such applications is the human eye. Conventional subjective assessment typically takes a four-step approach. First, pre-select a number of images from the input domain of a given image enhancement problem. Second, pick a set of competing enhancers, and generate the corresponding output images. Third, ask humans to rate the perceived quality of the enhanced images. Fourth, compare the enhancers according to the subjective results.

Unfortunately, conventional subjective assessment may suffer from three sources of biases. The first is the sampling bias. The underlying principle of conventional subjective assessment is to prove an enhancement method to be correct. This would require the set of pre-selected images to be large enough to sufficiently represent the input domain of interest. However, subject testing is an expensive and time-consuming endeavor. In practice, the number of images being examined is limited to a few hundreds (if not fewer), casting doubt on the assumption of sufficient sampling in the high-dimensional image space. The second is the *algorithmic bias*. It is important to note that the selection of test images precedes the selection of competing methods. One may take advantage of this (intentionally or unintentionally), and tunes her/his enhancer to overfit the pre-selected images, drawing overly optimistic conclusions

on the real-world generalization performance. The third is the *subjective bias*. That is, the test results may further be cherry-picked to bias towards certain methods. In summary, it is sad, but not uncommon, to see that a "state-of-the art" image enhancer produces superior results in its original paper, but remains particularly weak at handling examples appeared in subsequent work.

In this paper, we contribute to debiasing conventional subjective assessment by injecting an automated, adaptive and sample-efficient mechanism to select input domain samples. Our inspirations are drawn from interdisciplinary prior work on "model falsification as model comparison", a renowned philosophy in the fields of computational vision [72], software testing [43] and computer vision [47, 40, 62]. Specifically, we start from a large-scale image set as a finite approximation to the input space of an image enhancement application. According to the available human labelling budget, our method automatically selects a set of adaptive and diverse images for subsequent subjective testing. The selected images are optimal in terms of discriminating between the enhancers, while having the maximum withingroup variation in a latent space to ensure content diversity. Subjective results of the corresponding enhanced images reveal the advantages and disadvantages of the competing methods, and provide a debiased ranking of their relative performance. Our subjective assessment method is applicable to a wide variety of image processing and computer photography subfields, and we choose three real-world image enhancement applications as demonstration: 1) single image dehazing, 2) single image super-resolution, and 3) low-light image enhancement.

2. Related Work

We provide a concise overview of three real-world image enhancement applications, with emphasis on how previous subjective and objective assessments were carried out.

Single Image Dehazing Outdoor images are often captured in the presence of haze [46]. Due to the absorption and refraction of light by turbid medium, the resulting images may suffer from poor visibility and color shift. Conventional single image dehazing methods relied on the Koschmieder's model [18] and natural image priors [21, 23, 44, 87, 1] to estimate the atmospheric light and the transmission map. With the recent advances in convolutional neural networks (CNNs), plenty of CNN-based methods [77, 51, 32, 50, 2] have been proposed, directly regressing clean images from hazy ones.

Ma *et al.* [42] made initial attempts to subjective assessment of single image dehazing. A somewhat surprising observation is that due to the introduction of algorithmdependent distortions, the dehazed results by some algorithms are statistically insignificant compared to the input hazy images. Choi *et al.* [6] put emphasis on perceived fog density instead of overall quality. Li *et al.* [33] evaluated several dehazing algorithms for both human and machine vision. Tang *et al.* [58] investigated nighttime image dehazing, asking subjects to rate four aspects of dehazed algorithms: detail recognition, color fidelity, image authenticity, and overall effect.

The above subjective studies lead to an increasing consensus that objective quality models such as the mean squared error (MSE), the structural similarity (SSIM) index [71] and other no-reference methods [22] cannot accurately predict the perceived quality of dehazed images.

Single Image Super-Resolution Super-resolving a lowresolution image into a high-resolution one is very challenging, especially with a large scaling factor. Early attempts were mainly interpolation-based methods [80] using natural image statistics. In the late 2000s, model-based methods [11] came into play, with gradient profile prior [57], sparsity prior [75], and self-similarity prior [24] being representative. In the past five years, CNN-based methods began to dominate this field [10, 28, 29, 79, 84, 85], some of which were combined with generative adversarial networks (GANs) to encourage texture synthesis [86].

Yang et al. [74] presented one of the first subjective evaluations of single image super-resolution methods. Later, the authors [39] enlarged their dataset by including more highresolution images and more competing methods. Johnson et al. [27] performed a small-scale subjective experiment to verify the perceptual advantages of the VGG loss [55] in super-resolution. Gu et al. [15] established PIPAL - a largescale subject-rated dataset for image restoration, including GAN-based super-resolution results. With respect to objective assessment, MSE and SSIM [71] are still the most widely adopted metrics. LPIPS [82] and DISTS [9] that give credit to visually plausible synthesized textures have also been used for benchmarking purposes in some of the latest work. When high-resolution images are assumed unknown, NIQE is sometimes used for quantitative comparison.

Low-Light Image Enhancement Arguably the most significant impediment to high-quality pictures is lack of light [20]. On the one hand, a nighttime or indoor scene may not provide adequate light. On the other hand, a daytime scene may has a high dynamic range (*i.e.*, unbalanced lighting), causing current imaging techniques to collect insufficient light in shadow regions. Early computational methods for low-light image enhancement were equated to contrast enhancement either globally [7] or locally [56]. The Retinex theory [30] was also extensively studied in this context, where the problem of low-light enhancement is transformed to illumination map estimation [17]. Recently, many data-driven CNN-based methods [26, 73, 63, 26] with and with-



Figure 1: Overview of our debiased subjective assessment in the context of single image dehazing. (a): A large set of hazy images as an approximation to the input domain \mathcal{X} . (b): Top-K hazy images selected from (a) to best discriminate between Shao20 [53] and FFA-Net [48] by optimizing Eq. (2). D_1 and D_2 are implemented by DISTS [9] and MSE of the last feature layer of VGGNet [55], respectively. (c): Pairs of dehazed images corresponding to representative hazy images in (b).

out paired supervision have been developed, obtaining superior results on a limited number of visual examples.

Limited work has been done to assess low-light image enhancement subjectively. Hwang et al. [25] carried out a user study to validate their proposed enhancer using 20 low-contrast images, some of which are due to poor lighting conditions. Chen et al. [4] included images captured in hazy, underwater, and low-light conditions for human testing. A recent subjective study [16] compared six advanced low-light enhancers. Another small-scale subjective study was reported in [26] on 23 low-light images with six enhancement algorithms. With regard to objective assessment, MSE and SSIM prevail in this application. Using the input image as a corrupted reference, one may refer to VIF [54] and PCQI [64] for measuring the degree of enhancement. To the best of our knowledge, existing no-reference models [12] remain particularly weak at predicting the perceived quality of low-light enhanced images.

The above-mentioned subjective experiments may differ in how test images are presented to the subjects and how human data are collected, but they all need to pre-select the test images by the experimenters. Therefore, the results may suffer from sampling, algorithmic, and subjective biases, motivating us to debias subjective assessment of realworld image enhancement in this work.

3. Proposed Method

We formulate subjective assessment of real-world image enhancement in a general mathematical framework. Starting from an input image domain \mathcal{X} , we can easily sample an image $x \in \mathcal{X}$ that needs to be enhanced for improved visual quality. We choose a set of enhancement methods $\mathcal{F} = \{f_j\}_{j=1}^N$, each of which takes an $x \in \mathcal{X}$ as input, and produces an enhanced output $y_j = f_j(x)$. We also assume a subjective assessment environment, where human participants can reliably rate the perceived quality of y_j . The ultimate goal is to compare the N methods on the input domain \mathcal{X} containing enormous images, under the constraint of a very limited human labelling budget.

Conventional subjective assessment first pre-selects a small image set $S = \{x^{(i)}\}_{i=1}^M$. For each image $x \in S$, a set of enhanced versions $\{y_j\}_{j=1}^N$ are created, based on which subjective testing reveals the relative performance of $\{f_j\}_{j=1}^N$ on x. The model with the highest subjective ratings averaged over S is the best. As discussed previously, this would introduce several sources of biases. Inspired by interdisciplinary work under the scientific philosophy of "model falsification as model comparison" [43, 47, 40], especially following the well-established principle of maximum differentiation (MAD) competition [72], we aim to falsify an enhancer by finding a minimum set of images, which are most likely to be its counterexamples. An enhancer that is more difficult to be falsified is considered better.

We first describe the simplest situation, where two enhancers f_1 and f_2 are being compared, and the human labelling budget only allows us to select a single image $x \in \mathcal{X}$ for subjective testing. Then, the core question boils down to: How to automatically select which image for subjective testing from massive candidate images, such that the relative performance f_1 and f_2 may be most easily revealed?

According to the MAD competition methodology [72], our method selects the image $\hat{x} \in \mathcal{X}$ that best differentiates between f_1 and f_2 :

$$\hat{x} = \operatorname*{argmax}_{x \in \mathcal{X}} D_1(f_1(x), f_2(x)), \tag{1}$$

where D_1 is a quantitative measure to approximate the perceptual distance between $f_1(x)$ and $f_2(x)$. Visual inspection on $f_1(\hat{x})$ and $f_2(\hat{x})$ leads to two plausible results:

The majority of human subjects prefer f₁(x̂) (or f₂(x̂)) over f₂(x̂) (or f₁(x̂)). In this case, the proposed





(b) With D_2

Figure 2: Top-K images selected (a) without and (b) with the diversity loss, respectively.

subjective assessment method automatically detects a strong counterexample of one enhancer, not the other; a clear winner is declared. The chosen \hat{x} is the most informative in ranking the relative performance between f_1 and f_2 .

 Human subjects give f₁(x̂) and f₂(x̂) similar ratings. High rating indicates that both methods generate desirable but diverse outputs, which makes sense in real-world image enhancement that admits multiple plausible outputs. Low rating indicates that both fail, in dramatically different ways, to produce reasonable results. In either case, the chosen x̂ reveals different aspects of the strengths (or weaknesses) of f₁ and f₂, but contributes less to their relative performance ranking.

It seems straightforward to extend this idea to compare f_1 and f_2 on a small image subset $S \subset \mathcal{X}$ containing images with top-K largest distances computed by Eq. (1). However, such a naïve extension may simply identify algorithm failures of the same underlying root cause, leading to less interesting comparison (see Figure 2). To encourage more diverse failures of the competing models to be spotted, we modify Eq. (1) when looking for the k-th image:

$$\hat{x}^{(k)} = \underset{x \in \mathcal{X} \setminus \mathcal{S}}{\operatorname{argmax}} D_1(f_1(x), f_2(x)) + \lambda_1 D_2(x, \mathcal{S}), \quad (2)$$

where $S = {\hat{x}^{(i)}}_{i=1}^{k-1}$ is the set of k-1 images that have already been identified according to Eq. (2). D_2 is a second measure to quantify the semantic distance between an image x and the set S. λ_1 governs the trade-off between the two terms. Once $\hat{x}^{(k)}$ is identified, we incorporate it into S.

Given N enhancement algorithms, our subjective assessment method chooses top-K images for each of the $\binom{N}{2}$ distinct pairs of enhancers, gives rise to a final set \mathcal{D} with N(N-1)K enhanced images. It is worth noting that the size of \mathcal{D} is independent of the size of the input domain \mathcal{X} .

Algorithm 1: Debiased Subjective Assessment of Real-World Image Enhancement **Input:** A large-scale set \mathcal{X} , a list of competing methods $\mathcal{F} = \{f_j\}_{j=1}^N$, and two distance measures D_1 and D_2 **Output:** Global ranking vector $\mu \in \mathbb{R}^N$ 1 $\mathcal{D} \leftarrow \emptyset$ **2** for $j \leftarrow 1$ to N do 3 Compute the enhanced images $\{f_i(x)|x \in \mathcal{X}\}$ 4 end 5 for $i \leftarrow 1$ to N - 1 do for $j \leftarrow i + 1$ to N do 6 $\mathcal{S} \leftarrow \emptyset$ 7 for $k \leftarrow 1$ to K do 8 Select $\hat{x}^{(k)}$ by optimizing Eq. (2) 9 $\mathcal{S} \leftarrow \mathcal{S} \cup \hat{x}^{(k)}$ 10 $\mathcal{D} \leftarrow \mathcal{D} \cup \{f_i(\hat{x}^{(k)}), f_i(\hat{x}^{(k)})\}$ 11 end 12 end 13 14 end 15 Create the count matrix C for \mathcal{D} via the 2AFC method 16 Compute μ by optimizing Eq. (3)

Provided that the computational cost of image enhancement is negligible, the proposed subjective assessment method suggests expanding \mathcal{X} to cover as many images (and therefore diverse failures of the competing methods) as possible.

We now introduce the subjective assessment environment for gathering human opinions of image quality. As each input image $x \in S$ is associated with a pair of enhanced images $\{f_i(x), f_j(x)\} \subset D$, it is natural to employ the two-alternative forced choice (2AFC) method. That is, the subject is presented with $f_i(x)$ and $f_j(x)$ simultaneously, and is forced to choose which one is of higher quality. After subjective testing, we arrange the collected human data in an $N \times N$ matrix C, where C_{ij} records the number of votes for f_i and against f_j . Finally, we adopt maximum likelihood for multiple options [60] under the Thurstone's model [59] to infer the global ranking of \mathcal{F} . Specifically, we let μ be the vector of global ranking scores $[\mu_1, \mu_2, \ldots, \mu_N]$, and define the log-likelihood of the count matrix, C, as

$$L(\mu; C) = \sum_{ij} C_{ij} \log(\Phi(\mu_i - \mu_j)),$$
 (3)

where $\Phi(\cdot)$ is the standard Normal cumulative distribution function. When maximizing $L(\mu; C)$, one often adds an addition constraint, $\sum_i \mu_i = 0$, to obtain a unique optimal solution. We summarize the proposed debiased subjective assessment method in Algorithm 1, and show an overview



Figure 3: Global ranking results of (a) single image dehazing, (b) single image super-resolution, and (c) low-light image enhancement by optimizing Eq. (3).

of it in the context of single image dehazing in Figure 1.

4. Applications to Image Enhancement

In this section, we apply our subjective assessment method to three real-world image enhancement tasks: single image dehazing, single image super-resolution, and lowlight image enhancement.

4.1. Experimental Setups

Construction of \mathcal{X} For dehazing, the 10,000 real hazy images are originated from RESIDE [36] and the Internet. For super-resolution, the 10,000 low-resolution images are from WED [41], OST [67], and the Internet. For low-light enhancement, the 10,000 low-light images are chosen from ExDark [37], NPE [65], DICM [31], MBLLEN [38], VV [61], and the Internet. No manual pre-screening is needed at this stage.

Selection of \mathcal{F} For dehazing, we select eight popular algorithms published from 2015 to 2020: CAP [87], Berman16 [1], AOD-Net [32], Cho18 [5], GCANet [3], FFA-Net [48], Dhara20 [8], and Shao20 [53], among which CAP, Berman16, Cho18, and Dhara20 are knowledge-driven, while the rest are data-driven.

For super-resolution, we select eight CNN-based methods ranging from 2016 to 2020: SRCNN [10], EDSR [34], DBPN [19], TSRN [14], ESRGAN [68], RankSR-GAN [83], ESRGANplus [49], and USRGAN [78].

For low-light enhancement, we select eight methods from 2016 to 2020: Fu16 [13], BIMEF [76], Retinex-Net [73], JED [52], EnlightenGAN [26], Zhang19 [81], Wang19 [66], and Zero-DCE [16], among which Retinex-Net, EnlightenGAN, and Zero-DCE are CNN-based. The implementations of all 24 methods are obtained from the respective authors, and are tested with the default settings. **Construction of** S The created \mathcal{X} may be noisy, including images that lie out of the input domain of interest. Therefore, for dehazing, we replace a selected image that is either non-natural or haze-free with the next eligible one that optimizes Eq. (2). Moreover, for each $x \in S$, the visibility improvements in the corresponding "dehazed" images $f_i(x)$ and $f_j(x)$ are automatically checked by the computational method in [6]. If there is no predicted improvement in either dehazed image, we discard x. We apply the same image screening procedure for low-light enhancement, where the computational method in [12] is adopted for automatic verification of detail enhancement.

Subjective Experiment We conduct subjective user studies to gather human quality scores of the enhancement results in \mathcal{D} . The 2AFC method is adopted, allowing differentiation of subtler quality variations. Subjects are forced to choose the image with higher perceived quality with unlimited viewing time. For each enhancement application, we set K = 12, resulting in a total of $\binom{8}{2} \times 12 = 336$ paired comparisons. To relieve fatigue, subjects are allowed to take a break at anytime during each session of subjective testing. We gather data from 25 subjects with general background knowledge of image processing.

4.2. Main Results

Quantitative Results We show the global ranking results of the three real-world image enhancement applications in Figure 3, from which we have several interesting observations.

For *dehazing*, the main observation is that the syntheticto-real domain shift challenges all methods. Shao20 [53] leverages a bidirectional network to explicitly bridge the gap between the synthetic and realistic hazy images, and therefore exhibits the strongest generalization to the real



(f) GCANet/CAP (g) Shao20/Cho18 (h) CAP/GCANet (i) AOD-Net/Shao20 (j) FFA-Net/Berman16



world. By contrast, FFA-Net [48] relies exclusively on synthetic data for training, the majority of which are indoor scenes. Along with delicate feature attention and fusion modules, FFA-Net tends to overfit synthetic data, and has the worst performance in the debiased subjective experiment. Second, methods with less reliance on the Koschmieder's model [18] and image priors, such as GCANet [3] and Cho18 [5], generally outperform priorbased methods Berman16 [1], Dhara20 [8], and CAP [87], and the physical model-based AOD-Net [32]. This makes sense because current physical and statistical models oversimplify the natural imaging process in complex realistic hazy scenes, e.g., in the presence of non-uniform/dark light or heterogeneous haze density. As a result, algorithmdependent artifacts are likely to emerge, which may be perceptually more annoying than the haze (see Figure 4). Last, the Spearman's rank correlation coefficient (SRCC) between the subjective results of the competing methods and their publication times is only 0.167, implying the progress made in the field of single image dehazing might be somewhat over-estimated in terms of their real-world generalization, despite outstanding (synthetic) benchmark numbers.

For *super-resolution*, steady progress over the years has been reported in our experiment, with an SRCC value of 0.958 between subjective results and published years. SR-CNN [10] is the first CNN for super-resolution with three convolutions, and can be viewed as an end-to-end trainable sparse-coding based method [75]. EDSR [34] adds more convolution layers with residual connections to stabilize training. DBPN [19] replaces single-stage upsampling with iterative up/downsampling. TSRN [14] optimizes for the texture-aware LPIPS [82] metric, and underperforms RankSRGAN [83], ESRGAN [68], and ESRGANplus [49] based on GANs with stronger texture synthesis capabilities. The lastest USRGAN[78] inherits the flexibility of modelbased methods, while maintaining the end-to-end training capability of learning-based methods.

For low-light enhancement, the main observation is that CNN-based methods have not come to dominate this field due to the lack of ground-truth normal-light images for paired supervision. For example, Retinex-Net [73], ranked in the last place, only sees 485 realistic pairs during training, which may be insufficient to cover the real-world scene complexities. Zero-DCE [16] optimizes a CNN for a combination of image naturalness measures, including spatial consistency, exposedness, color constancy, and illumination smoothness, without reference to normal-light images. However, the combined loss has not be calibrated against human judgments, leading to unpredictable real-world generalization. An exception is EnlightenGAN [26], which leverages an unsupervised GAN to regularize the unpaired training, leading to the best performance in our subjective study. Second, it is difficult to enhance the details of lowlight images without amplifying the background noise. JED [52] performs joint optimization of low-light enhancement and noise suppression, leading to noticeable visual quality improvements. The multi-exposure fusion framework adopted by BIMEF [76] and Zhang19 [81] achieves a similar effect of noise reduction with comparable performance. Third, relying on classic image processing techniques such as multi-scale decomposition and adaptive histogram equalization, Fu16 [13] and Wang19 [66] tend to overshoot local details at the sacrifice of global brightness and contrast. Last, similar as dehazing, steady progress over the years is not reflected in our debiased subjective experiment, with an SRCC value of 0.071 between the subjective results and the



(a) USRGAN/ESRGAN

(e) TSRN/RankSRGAN



Figure 5: Representative distortions created by super-resolution methods in our experiment. f_i/f_j below each image means that f_i is used to produce the image, while f_j is the paired method for selecting the corresponding low-resolution image.



(f) Retinex-Net/Fu16

(g) EnlightenGAN/Fu16

(h) JED/Zero-DCE

(j) Zero-DCE/Zhang19

Figure 6: Representative distortions created by low-light enhancers in our experiment. f_i/f_j below each image means that f_i is used to produce the image, while f_j is the paired method for selecting the corresponding low-light image.

published years of the competing methods.

Qualitative Results We show some visual examples for each of the three tasks, summarizing and diagnosing the identified distortion patterns.

For dehazing in Figure 4, the perceived distortions can be approximately classified into five types: JPGE blocking (see (a), (b), (c), and (d)), color cast (see (b), (d), and (e)), loss of high-frequency information (see (f), and (g)), low-brightness (see (h) and (i)), and haze (see (i) and (j)). Knowledge-driven methods such as Dhara20 [8] and Cho18 derlying JPEG artifacts of hazy images from the Internet accompanied by the color cast problem. Data-driven methods such as AOD-net [32] and FFA-Net [48] have learned a more conservative dehazing strategy. CAP [87] tends to regions darker and the hazy regions hazier. Despite the best performance, Shao20 [53] is likely to smooth highfrequency details, which is successfully spotted by our debiased subjective method in (g).

For super-resolution in Figure 5, the perceived distor-(j)), fake textures (see (b), (c), (d), and (e)), incorrect se-(see (a) and (b)). CNNs not optimized for texture-aware

Method	Global Ranking		
	DISTS	Δ MSE	Δ SSIM
USRGAN [78]	1	0	0
ESRGANplus [49]	2	0	0
ESRGAN [68]	3	-1	0
RankSRGAN [83]	4	+1	0
TSRN [14]	5	0	-1
DBPN [19]	6	0	+1
EDSR [34]	7	0	0
SRCNN [10]	8	0	0

losses often suffer from blurring artifacts. CNNs reinforced by GANs are capable of synthesizing random textures, but remain weak at super-resolving structured (especially periodic) textures. All methods fail when it comes to images with rich semantics such as faces, validating face hallucination [35] as a separate super-resolution problem of its unique challenge and independent interest. With more specialized modules proposed, the field of single image superresolution begins to enter the era of local contrast overenhancement, as pointed out by our subjective method.

For *low-light enhancement* in Figure 6, the perceived distortions roughly belong to five classes: noise (see (a), (b), and (c)), JPEG blocking (see (d)), abnormal brightness (see (e) and (f)), color cast (see (g) and (h)), and poor exposure (see (i) and (j)). Similar as dehazing, knowledge-driven methods (*e.g.*, Fu16 [13], Wang19 [66], and Zhang19 [81]) encourage over-enhancing details, which significantly amplifies background noise and possible JPEG blocking. Unlike super-resolution, data-driven methods (*e.g.*, Retinex-Net [73] and Zero-DCE [16]) are far more brittle than knowledge-driven ones, which sometimes have unexpected behaviors, producing results with unnatural appearances. The best performer EnlightenGAN [26] exhibits the least amount of artifacts, but still appears to have halos around light sources in the scene, which is identified by our method.

Ablation Study We first analyze the sensitivity of our subjective results to different distance measures D_1 in Eq. (2). We use another two widely adopted metrics in signal and image processing - MSE and SSIM [71]. We opt for single image super-resolution, and follow the procedure in Section 4.1 to sample two subsets, each of which contains 336 pairs of images. We gather human data from 21 subjects. Table 1 shows the results, where we find that the global ranking is consistent across the three metrics. This may be because MAD chooses images to optimally discriminate between two models with large perceptual distances, which can be well approximated by all the three measures.

Next, we analyze the sensitivity of the obtained results



Figure 7: The SRCC values between the top-12 and other top-K rankings, where $K \in \{1, 2, ..., 11\}$.

5. Conclusion

Acknowledgments

References

- D. Berman, T. Treibitz, and S. Avidan. Non-local image dehazing. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1674–1682, 2016.
- [2] B. Cai, X. Xu, K. Jia, C. Qing, and D. Tao. DehazeNet: An end-to-end system for single image haze removal. *IEEE Transactions on Image Processing*, 25(11):5187–5198, 2016.
- [3] D. Chen, M. He, Q. Fan, J. Liao, L. Zhang, D. Hou, L. Yuan, and G. Hua. Gated context aggregation network for image dehazing and deraining. In *IEEE Winter Conference on Applications of Computer Vision*, pages 1375–1383, 2019.
- [4] Z. Chen, T. Jiang, and Y. Tian. Quality assessment for comparing image enhancement algorithms. In *IEEE Conference* on Computer Vision and Pattern Recognition, pages 3003– 3010, 2014.
- [5] Y. Cho, J. Jeong, and A. Kim. Model assisted multi-band fusion for single image enhancement and applications to robot vision. *IEEE Robotics and Automation Letters*, 3(4):2822– 2829, 2018.
- [6] L. K. Choi, J. You, and A. C. Bovik. Referenceless prediction of perceptual fog density and perceptual image defogging. *IEEE Transactions on Image Processing*, 24(11):3888–3901, 2015.
- [7] D. Coltuc, P. Bolon, and J. M. Chassery. Exact histogram specification. *IEEE Transactions on Image Processing*, 15(5):1143–1152, 2006.
- [8] S. K. Dhara, M. Roy, D. Sen, and P. K. Biswas. Color cast dependent image dehazing via adaptive airlight refinement and non-linear color balancing. *IEEE Transactions on Circuits* and Systems for Video Technology, to appear, 2020.
- [9] K. Ding, K. Ma, S. Wang, and E. P. Simoncelli. Image quality assessment: Unifying structure and texture similarity. arXiv preprint:2004.07728, 2020.
- [10] C. Dong, C. C. Loy, K. He, and X. Tang. Image super-resolution using deep convolutional networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(2):295–307, 2016.
- [11] W. Dong, L. Zhang, G. Shi, and X. Wu. Image deblurring and super-resolution by adaptive sparse domain selection and adaptive regularization. *IEEE Transactions on Image Processing*, 20(7):1838–1857, 2011.
- [13] X. Fu, D. Zeng, Y. Huang, Y. Liao, X. Ding, and J. Paisley. A fusion-based enhancing method for weakly illuminated images. *Information Sciences*, 129(12):82–96, 2016.
- [14] M. W. Gondal, B. Schölkopf, and M. Hirsch. The unreasonable effectiveness of texture transfer for single image superresolution. In *European Conference on Computer Vision*, pages 80–97, 2019.

- [15] J. Gu, H. Cai, H. Chen, X. Ye, J. Ren, and C. Dong. PIPAL: A large-scale image quality assessment dataset for perceptual image restoration. In *European Conference on Computer Vision*, pages 633–651, 2020.
- [16] C. Guo, C. Li, J. Guo, C. C. Loy, J. Hou, S. Kwong, and R. Cong. Zero-reference deep curve estimation for low-light image enhancement. In *IEEE Conference on Computer Vi*sion and Pattern Recognition, pages 1780–1789, 2020.
- [17] X. Guo, Y. Li, and H. Ling. Lime: Low-light image enhancement via illumination map estimation. *IEEE Transactions on Image Processing*, 26(2):982–993, 2016.
- [18] K. Harald. Theorie der horizontalen sichtweite: Kontrast und sichtweite. *Keim & Nemnich, Munich*, 12, 1924.
- [19] M. Haris, G. Shakhnarovich, and N. Ukita. Deep backprojection networks for super-resolution. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1664–1673, 2018.
- [20] S. W. Hasinoff, D. Sharlet, R. Geiss, A. Adams, J. T. Barron, F. Kainz, J. Chen, and M. Levoy. Burst photography for high dynamic range and low-light imaging on mobile cameras. *ACM Transactions on Graphics*, 35(6):1–12, 2016.
- [21] N. Hautière, J. P. Tarel, and D. Aubert. Towards fog-free in-vehicle vision systems through contrast restoration. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2007.
- [22] N. Hautière, J. P. Tarel, D. Aubert, and E. Dumont. Blind contrast enhancement assessment by gradient ratioing at visible edges. *Image Analysis & Stereology*, 27(2):87–95, 2008.
- [23] K. He, J. Sun, and X. Tang. Single image haze removal using dark channel prior. *IEEE Transactions on Pattern Analysis* and Machine Intelligence, 30(12):2341–2353, 2010.
- [24] J. B. Huang, A. Singh, and N. Ahuja. Single image superresolution from transformed self-exemplars. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 5197–5206, 2015.
- [25] S. J. Hwang, A. Kapoor, and S. B. Kang. Context-based automatic local image enhancement. In *European Conference* on Computer Vision, pages 569–582, 2012.
- [26] Y. Jiang, X. Gong, D. Liu, Y. Cheng, C. Fang, X. Shen, J. Yang, P. Zhou, and Z. Wang. EnlightenGAN: Deep light enhancement without paired supervision. arXiv preprint:1906.06972, 2019.
- [27] J. Johnson, A. Alahi, and F. Li. Perceptual losses for realtime style transfer and super-resolution. In *European Conference on Computer Vision*, pages 694–711, 2016.
- [28] J. Kim, J. Kwon Lee, and K. Mu Lee. Accurate image superresolution using very deep convolutional networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1646–1654, 2016.
- [29] W. S. Lai, J. B. Huang, N. Ahuja, and M. H. Yang. Deep Laplacian pyramid networks for fast and accurate superresolution. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 5835–5843, 2017.

- [30] E. H. Land. The retinex theory of color vision. *Scientific American*, 237(6):108–129, 1977.
- [31] C. Lee, C. Lee, and C. S. Kim. Contrast enhancement based on layered difference representation. In *IEEE International Conference on Image Processing*, pages 965–968, 2012.
- [32] B. Li, X. Peng, Z. Wang, J. Xu, and D. Feng. Aod-Net: Allin-one dehazing network. In *IEEE international Conference* on Computer Vision, pages 4770–4778, 2017.
- [33] B. Li, W. Ren, D. Fu, D. Tao, D. Feng, W. Zeng, and Z. Wang. Benchmarking single-image dehazing and beyond. *IEEE Transactions on Image Processing*, 28(1):492– 505, 2019.
- [34] B. Lim, S. Son, H. Kim, S. Nah, and K. Mu Lee. Enhanced deep residual networks for single image super-resolution. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 1132–1140, 2017.
- [35] C. Liu, H. Y. Shum, and W. T. Freeman. Face hallucination: Theory and practice. *International Journal of Computer Vision*, 75(1):115–134, 2007.
- [36] Y. Liu, G. Zhao, B. Gong, Y. Li, R. Raj, N. Goel, S. Kesav, S. Gottimukkala, Z. Wang, and W. Ren. Improved techniques for learning to dehaze and beyond: A Collective Studys. *arXiv preprint: 1807.00202*, 2018.
- [37] Y. P. Loh and C. S. Chan. Getting to know low-light images with the exclusively dark dataset. *Computer Vision and Image Understanding*, 178(3):30–42, 2019.
- [38] F. Lv, F. Lu, J. Wu, and C. Lim. MBLLEN: Low-light image/video enhancement using cnns. In *British Machine Vi*sion Conference, pages 1–13, 2018.
- [39] C. Ma, C. Y. Yang, X. Yang, and M. H. Yang. Learning a noreference quality metric for single-image super-resolution. *Computer Vision and Image Understanding*, 158(1):1–16, 2017.
- [40] K. Ma, Z. Duanmu, Z. Wang, Q. Wu, W. Liu, H. Yong, H. Li, and L. Zhang. Group maximum differentiation competition: Model comparison with few samples. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(4):851–864, 2020.
- [41] K. Ma, Z. Duanmu, Q. Wu, Z. Wang, H. Yong, H. Li, and L. Zhang. Waterloo Exploration Database: New challenges for image quality assessment models. *IEEE Transactions on Image Processing*, 26(2):1004–1016, 2016.
- [42] K. Ma, W. Liu, and Z. Wang. Perceptual evaluation of single image dehazing algorithms. In *IEEE International Conference on Image Processing*, pages 3600–3604, 2015.
- [43] W. M. McKeeman. Differential testing for software. *Digital Technical Journal*, 10(1):100–107, 1998.
- [44] G. Meng, Y. Wang, J. Duan, S. Xiang, and C. Pan. Efficient image dehazing with boundary constraint and contextual regularization. In *IEEE Conference on Computer Vision* and Pattern Recognition, pages 617–624, 2013.
- [45] A. Mittal, R. Soundararajan, and A. C. Bovik. Making a "completely blind" image quality analyzer. *IEEE Signal Pro*cessing Letters, 20(3):209–212, 2012.

- [46] S. G. Narasimhan and S. K. Nayar. Vision and the atmosphere. *International Journal of Computer Vision*, 48(3):233–254, 2002.
- [47] K. Pei, Y. Cao, J. Yang, and S. Jana. DeepXplore: Automated whitebox testing of deep learning systems. In *Symposium on Operating Systems Principles*, pages 1–18, 2017.
- [48] X. Qin, Z. Wang, Y. Bai, X. Xie, and H. Jia. FFA-Net: Feature fusion attention network for single image dehazing. In AAAI Conference on Artificial Intelligence, pages 11908– 11915, 2020.
- [49] N. C. Rakotonirina and A. Rasoanaivo. ESRGAN+ : Further improving enhanced super-resolution generative adversarial network. In *IEEE International Conference on Acoustics*, *Speech and Signal Processing*, pages 3637–3641, 2020.
- [50] W. Ren, S. Liu, H. Zhang, J. Pan, X. Cao, and M. H. Yang. Single image dehazing via multi-scale convolutional neural networks. In *European Conference on Computer Vision*, pages 154–169, 2016.
- [51] W. Ren, L. Ma, J. Zhang, J. Pan, X. Cao, W. Liu, and M.-H. Yang. Gated fusion network for single image dehazing. In *IEEE Conference on Computer Vision and Pattern Recogni*tion, pages 3253–3261, 2018.
- [52] X. Ren, M. Li, W. H. Cheng, and J. Liu. Joint enhancement and denoising method via sequential decomposition. In *IEEE International Symposium on Circuits and Systems*, pages 1– 5, 2018.
- [53] Y. Shao, L. Li, W. Ren, C. Gao, and N. Sang. Domain adaptation for image dehazing. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2808–2817, 2020.
- [54] H. R. Sheikh and A. C. Bovik. Image information and visual quality. *IEEE Transactions on Image Processing*, 15(2):430– 444, 2006.
- [55] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, pages 1–14, 2015.
- [56] J. A. Stark. Adaptive image contrast enhancement using generalizations of histogram equalization. *IEEE Transactions on Image Processing*, 9(5):889–896, 2000.
- [57] J. Sun, Z. Xu, and H. Y. Shum. Image super-resolution using gradient profile prior. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008.
- [58] Q. Tang, J. Yang, X. He, W. Jia, Q. Zhang, and H. Liu. Nighttime image dehazing based on retinex and dark channel prior using Taylor series expansion. *Computer Vision and Image Understanding*, 202(2):103086, 2021.
- [59] L. L. Thurstone. A Law of Comparative Judgment. *Psychological review*, 34(4):273, 1927.
- [60] K. Tsukida and M. R. Gupta. How to analyze paired comparison data. Technical report, Washington University Seattle Department of Electrical Engineering, 2011.
- [61] V. Vonikakis, R. Kouskouridas, and A. Gasteratos. On the evaluation of illumination compensation algorithms. *Multimedia Tools and Applications*, 77(8):9211–9231, 2018.

- [62] H. Wang, T. Chen, Z. Wang, and K. Ma. I am going MAD: Maximum discrepancy competition for comparing classifiers adaptively. *arXiv preprint:2002.10648*, 2020.
- [63] R. Wang, Q. Zhang, C. W. Fu, X. Shen, W. S. Zheng, and J. Jia. Underexposed photo enhancement using deep illumination estimation. In *IEEE Conference on Computer Vision* and Pattern Recognition, pages 6849–6857, 2019.
- [64] S. Wang, K. Ma, H. Yeganeh, Z. Wang, and W. Lin. A patch-structure representation method for quality assessment of contrast changed images. *IEEE Signal Processing Letters*, 22(12):2387–2390, 2015.
- [65] S. Wang, J. Zheng, H. M. Hu, and B. Li. Naturalness preserved enhancement algorithm for non-uniform illumination images. *IEEE Transactions on Image Processing*, 22(9):3538–3548, 2013.
- [66] W. Wang, Z. Chen, X. Yuan, and X. Wu. Adaptive image enhancement method for correcting low-illumination images. *Information Sciences*, 496(2):25–41, 2019.
- [67] X. Wang, K. Yu, C. Dong, and C. C. Loy. Recovering realistic texture in image super-resolution by deep spatial feature transform. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 606–615, 2018.
- [68] X. Wang, K. Yu, S. Wu, J. Gu, Y. Liu, C. Dong, Y. Qiao, and C. C. Loy. ESRGAN: Enhanced super-resolution generative adversarial networks. In *European Conference on Computer Vision*, pages 63–79, 2019.
- [69] Z. Wang and A. C. Bovik. Modern Image Quality Assessment. Synthesis Lectures on Image, Video, and Multimedia Processing, 2(1):1–156, 2006.
- [70] Z. Wang and A. C. Bovik. Reduced-and no-reference image quality assessment. *IEEE Signal Processing Magazine*, 28(6):29–40, 2011.
- [71] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
- [72] Z. Wang and E. P. Simoncelli. Maximum Differentiation (MAD) Competition: A methodology for comparing computational models of perceptual quantities. *Journal of Vision*, 8(12):1–13, 2008.
- [73] C. Wei, W. Wang, W. Yang, and J. Liu. Deep retinex decomposition for low-light enhancement. In *British Machine Vision Conference*, pages 127–136, 2018.
- [74] C. Y. Yang, C. Ma, and M. H. Yang. Single-image superresolution: A benchmark. In *European Conference on Computer Vision*, pages 372–386, 2014.
- [75] J. Yang, J. Wright, T. S. Huang, and Y. Ma. Image superresolution via sparse representation. *IEEE Transactions on Image Processing*, 19(11):2861–2873, 2010.
- [76] Z. Ying, G. Li, and W. Gao. A bio-inspired multi-exposure fusion framework for low-light image enhancement. arXiv preprint:1711.00591, 2017.

- [77] H. Zhang and V. M. Patel. Densely connected pyramid dehazing network. In *IEEE Conference on Computer Vision* and Pattern Recognition, pages 3194–3203, 2018.
- [78] K. Zhang, L. Van Gool, and R. Timofte. Deep unfolding network for image super-resolution. In *IEEE Conference* on Computer Vision and Pattern Recognition, pages 3217– 3226, 2020.
- [79] K. Zhang, W. Zuo, and L. Zhang. Learning a single convolutional super-resolution network for multiple degradations. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3262–3271, 2018.
- [80] L. Zhang and X. Wu. An edge-guided image interpolation algorithm via directional filtering and data fusion. *IEEE Transactions on Image Processing*, 15(8):2226–2238, 2006.
- [81] Q. Zhang, Y. Nie, and W.-S. Zheng. Dual illumination estimation for robust exposure correction. *Computer Graphics Forum*, 38(7):243–252, 2019.
- [82] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 586–595, 2018.
- [83] W. Zhang, Y. Liu, C. Dong, and Y. Qiao. RankSRGAN: Generative adversarial networks ranker for image superresolution. In *IEEE International Conference on Computer Vision*, pages 3096–3105, 2019.
- [84] Y. Zhang, K. Li, K. Li, L. Wang, B. Zhong, and Y. Fu. Image super-resolution using very deep residual channel attention networks. In *European Conference on Computer Vision*, pages 294–310, 2018.
- [85] Y. Zhang, Y. Tian, Y. Kong, B. Zhong, and Y. Fu. Residual dense network for image super-resolution. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2472–2481, 2018.
- [86] J. Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired imageto-image translation using cycle-consistent adversarial networks. In *IEEE International Conference on Computer Vi*sion, pages 2223–2232, 2017.
- [87] Q. Zhu, J. Mai, and L. Shao. A fast single image haze removal algorithm using color attenuation prior. *IEEE Transactions on Image Processing*, 24(11):3522–3533, 2015.