
Comparison of Image Quality Models for Optimization of Image Processing Systems

Keyan Ding¹ · Kede Ma¹ · Shiqi Wang¹ · Eero P. Simoncelli²

Abstract The performance of objective image quality assessment (IQA) models has been evaluated primarily by comparing model predictions to human judgments. Perceptual datasets (e.g., LIVE and TID2013) gathered for this purpose provide useful benchmarks for improving IQA methods, but their heavy use creates a risk of overfitting. Here, we perform a large-scale comparison of perceptual IQA models in terms of their use as objectives for the optimization of image processing algorithms. Specifically, we evaluate eleven full-reference IQA models by using them as objective functions to train deep neural networks for four low-level vision tasks: denoising, deblurring, super-resolution, and compression. Extensive subjective testing on the optimized images allows us to rank the competing models in terms of their perceptual performance, elucidate their relative advantages and disadvantages for these tasks, and propose a set of desirable properties for incorporation into future IQA models.

Keywords Image quality assessment · Perceptual optimization · Performance evaluation

Keyan Ding
E-mail: keyan.ding@my.cityu.edu.hk

Kede Ma
E-mail: kede.ma@cityu.edu.hk

Shiqi Wang
E-mail: shiqwang@cityu.edu.hk

Eero P. Simoncelli
E-mail: eero.simoncelli@nyu.edu

¹ Department of Computer Science, City University of Hong Kong, Kowloon, Hong Kong

² Howard Hughes Medical Institute, Center for Neural Science, and Courant Institute of Mathematical Sciences, New York University, New York, USA

1 Introduction

The goal of objective image quality assessment (IQA) is the construction of computational models that predict the perceived quality of visual images. Such models can be used to evaluate and compare image processing methods and systems (Wang, 2011). The standard paradigm for testing IQA models is to compare them to human perceptual quality ratings of distorted images, many of which are available in datasets such as LIVE (Sheikh et al., 2006) or TID2013 (Ponomarenko et al., 2015). However, excessive reuse of these test sets during IQA model development may lead to overfitting, and as a consequence, poor generalization to images corrupted by distortions that are not in the test sets (see Table 3).

A highly promising (but relatively under-studied) application of objective IQA measures is to guide the design and optimization of new image processing algorithms. The parameters of image processing methods are usually adjusted to minimize the mean squared error (MSE), the simplest of all fidelity metrics, despite the fact that it has been widely criticised for its poor correlation with human perception of image quality (Girod, 1993). Early attempts at perceptual optimization used the *structural similarity* (SSIM) index (Wang et al., 2004) in place of MSE, and achieved modest gains in applications of image restoration (Channappayya et al., 2008), wireless video streaming (Vukadinovic and Karlsson, 2009), video coding (Wang et al., 2011), and image synthesis (Snell et al., 2017). Driven by the surge of impressive results in deep neural networks (DNNs), recent authors have used perceptual measures based on pre-trained DNNs for optimization purposes (Johnson et al., 2016), although these have not been tested against human judgments.

In this paper, we systematically evaluate a large set of full-reference IQA models in the context of perceptual optimization. To determine their suitability for optimization, we first test the models on recovering a reference image from a given initialization by optimizing the model-reported distance to the reference. For many IQA methods, we find that the optimization does not converge to the reference image, and often creates severe distortions. These optima are either local (due to non-convexity of the IQA objective), or global but non-unique (due to the loss of information in the IQA objective). We select a subset of eleven IQA models as useful “perceptual” objectives for use in optimizing DNNs in four low-level vision tasks - image denoising, blind image deblurring, single image super-resolution, and lossy image compression. Extensive subjective testing on the optimized images reveals the relative performance of the competing models. Moreover, careful inspection of their visual failures indicates limitations in model design, which in turn sheds light on developing future IQA models in a principled way.

2 Taxonomy of Full-Reference IQA Models

Full-reference IQA methods can be broadly classified into five categories:

- *Error visibility methods.* These apply a distance measure directly to pixels (such as MSE), or to transformed representations of the images. The MSE possesses useful properties for optimization (e.g., differentiability and convexity), and when combined with linear-algebraic tools, analytical solutions can often be obtained. For example, the classical solution to the MSE-optimal denoising problem (assuming a translation-invariant Gaussian signal model) is the Wiener filter (Weiner, 1950). Given that MSE in the pixel domain is poorly correlated with perceived image quality, many IQA models operate by first mapping to a more perceptually appropriate representation (Safranek and Johnston, 1989; Daly, 1992; Lubin, 1993; Watson, 1993; Teo and Heeger, 1994; Watson et al., 1997; Larson and Chandler, 2010; Laparra et al., 2016).
- *Structural similarity (SSIM) methods.* These are constructed to measure the similarity of local image “structures”. The prototype is the SSIM index (Wang et al., 2004), which combines similarity measures of three conceptually independent components - luminance, contrast and structure. It has become a *de facto* standard in the field of perceptual image processing, and has inspired subsequent IQA models based on feature similarity (Zhang et al., 2011), gradient similarity (Liu et al., 2012a), edge strength similarity (Zhang et al., 2013), and saliency similarity (Zhang et al., 2014).
- *Information theoretic methods.* These attempt to measure some approximation of the mutual information between the perceived reference and distorted images as an indication of perceptual image quality. Statistical modeling of the image source, the distortion process, and the HVS is critical in algorithm development. The prototype is the visual information fidelity (VIF) measure (Sheikh and Bovik, 2006).
- *Learning-based methods.* These learn the relationship between the input images and the perceptual distance from a large set of examples, using supervised machine learning methods. By leveraging the power of DNNs, these methods have come to dominate the field of IQA, in terms of performance on existing image quality databases (Bosse et al., 2018; Prashnani et al., 2018). But given the high dimensionality of the input space (i.e., the number of pixels, typically millions), these methods are prone to overfitting the data. Strategies that compensate for the insufficiency of labeled training data include building on pre-trained networks (Zhang et al., 2018; Ding et al., 2020), training on local image patches (Bosse et al., 2018), and combining multiple IQA databases (Zhang et al., 2019b).
- *Fusion-based methods.* These aim to combine existing IQA methods to build a “super-evaluator” that exploits the diversity and complementarity of the incorporated methods for improved quality prediction performance (analogous to “boosting” methods in machine learning). Fusion combinations can be determined empirically (Ye et al., 2014) or learned from data (Liu et al., 2012b; Ma et al., 2019). Some methods incorporate deterministic or statistical image priors to regularize an IQA measure (Jordan, 1881; Ulyanov et al., 2018). Since such regularizers can be seen as a form of no-reference IQA measures (Wang and Bovik, 2011), we also view these as fusion solutions.

3 Screening of Full-Reference IQA Models for Perceptual Optimization

We used a naïve task to demonstrate the issues encountered when using IQA models in gradient-based perceptual optimization. This task also allows us to pre-screen existing models, and to motivate the design of experiments used in subsequent comparisons.

3.1 Reference Image Recovery

Given a reference (undistorted) image x and an initial image y_0 , we aim to recover x by numerically optimizing

$$y^* = \arg \min_y D(x, y), \quad (1)$$

where D denotes a full-reference IQA measure with a lower score indicating higher predicted quality, and y^* is the recovered image. As an example, if D is the MSE, the (trivial) analytical solution is $y^* = x$, indicating full recoverability. For the majority of current IQA models, which are continuous and differentiable, solutions must be sought numerically, using gradient-based iterative solvers. We consider a total of 17 methods: three error visibility methods - MAD (Larson and Chandler, 2010), PAMSE (Xue et al., 2013) and NLPD (Laparra et al., 2016), seven structural similarity methods - MS-SSIM (Wang et al., 2003), CW-SSIM (Wang and Simoncelli, 2005), FSIM (Zhang et al., 2011), SFF (Chang et al., 2013), GMSD (Xue et al., 2014) and VSI (Zhang et al., 2014), MCSD (Wang et al., 2016), two information theoretical methods - IFC (Sheikh et al., 2005) and VIF (Sheikh and Bovik, 2006), and five deep learning-based methods - GTI-CNN (Ma et al., 2018), DeepIQA (Bosse et al., 2018), PieAPP (Prashnani et al., 2018), LPIPS (Zhang et al., 2018) and DISTs (Ding et al., 2020). As this paper focuses on the perceptual optimization performance of individual IQA measures, fusion-based methods are not included.

Fig. 1 and Fig. 2 show the recovery results from two different initializations - a white Gaussian noise image and a JPEG compressed version of the reference image. For all methods, we find that optimization converges to a final image with a substantially lower score than that of the initial image. All models based on injective mappings (MS-SSIM, PAMSE, NLPD and DISTs) are guaranteed to recover the reference image (although the convergence may depend on the choice of initial image). Many of the remaining IQA models fail to recover the main structures of the reference image when initialized with the white Gaussian noise image, or create noticeable model-dependent distortions when initialized with the JPEG compressed image. This is because these methods rely on surjective mapping functions to transform the images to a “perceptual” space for quality computation. For example, GTI-CNN (Ma et al., 2018) uses a surjective DNN with four stages of convolution, subsampling, and halfwave rectification. The resulting undercomplete representation is optimized for geometric transformation invariance, at the cost of significant information loss. The examples demonstrate that preservation of some aspects

of this lost information is important for perceptual quality. Similar arguments can be applied to other surjective DNN-based IQA models, such as DeepIQA (Bosse et al., 2018) and PieAPP (Prashnani et al., 2018). In addition, with a better initialization (e.g., a JPEG compressed image with roughly correct local luminances), optimization guided by surjective models achieves a perceptually better image, compared to initialization with purely white Gaussian noise. Nevertheless, the visual quality of the final images is in some cases worse than that of the initial JPEG image (see Fig. 2 (h), (n) (o) and (p)).

3.2 IQA Model Selection

The reference image recovery test results were used to pre-screen the full set of IQA models, excluding those that perform poorly (due to surjectivity) or closely (due to similar design). The following 11 full-reference IQA models were selected for subsequent evaluation:

1. MAE, the Mean Absolute Error (ℓ_1 -norm) of pixel values, has been frequently adopted in optimization similar as MSE, despite its poor perceptual relevance. MAE has been shown to consistently outperform MSE in image restoration tasks (Zhao et al., 2016).
2. MS-SSIM (Wang et al., 2003), the Multi-Scale extension of the SSIM index (Wang et al., 2004), provides more exibility than single-scale SSIM, allowing for a wider range of viewing distances. MS-SSIM has become a standard “perceptual” quality measure, and has guided the design of DNN-based image super-resolution (Zhao et al., 2016; Snell et al., 2017) and compression (Ballé et al., 2018) algorithms.
3. VIF (Sheikh and Bovik, 2006), the Visual Information Fidelity measure, predicts the quality of the distorted image by quantifying how much information in the reference image is preserved. VIF can be computed either in spatial or wavelet (Simoncelli et al., 1992) domain. A distinct property of VIF relative to other IQA models is that it can indicate that the “distorted” image has visual quality superior to that of the reference (Wang et al., 2015).
4. CW-SSIM (Wang and Simoncelli, 2005), the Complex Wavelet SSIM index, is designed to be robust to small geometric distortions such as translation and rotation. The construction allows for consistent phase shifts of wavelet coefficients, which preserves local image features. CW-SSIM addresses the limitation of most IQA methods that require a precise registration process at the front end.

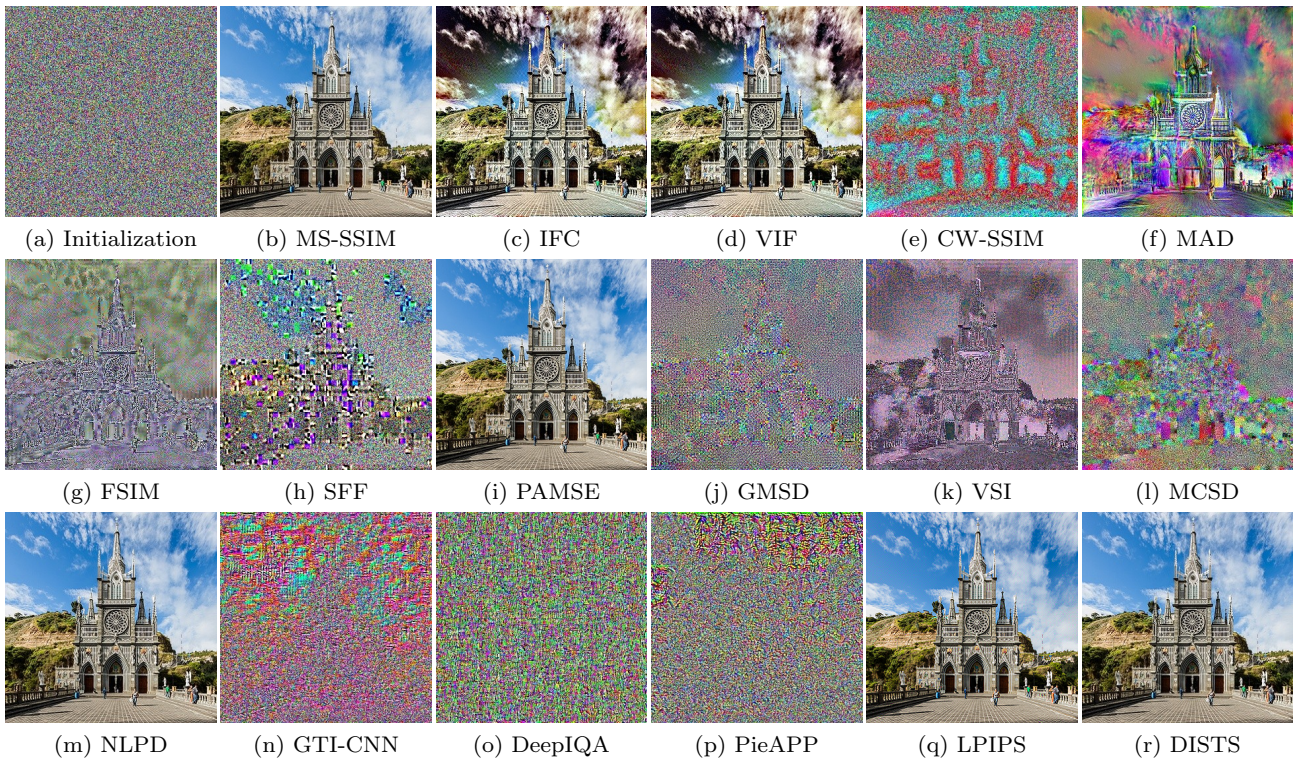


Fig. 1 Reference image recovery results, initialized from (a) a white Gaussian noise image. (b)-(r) Images recovered by optimizing different IQA models.

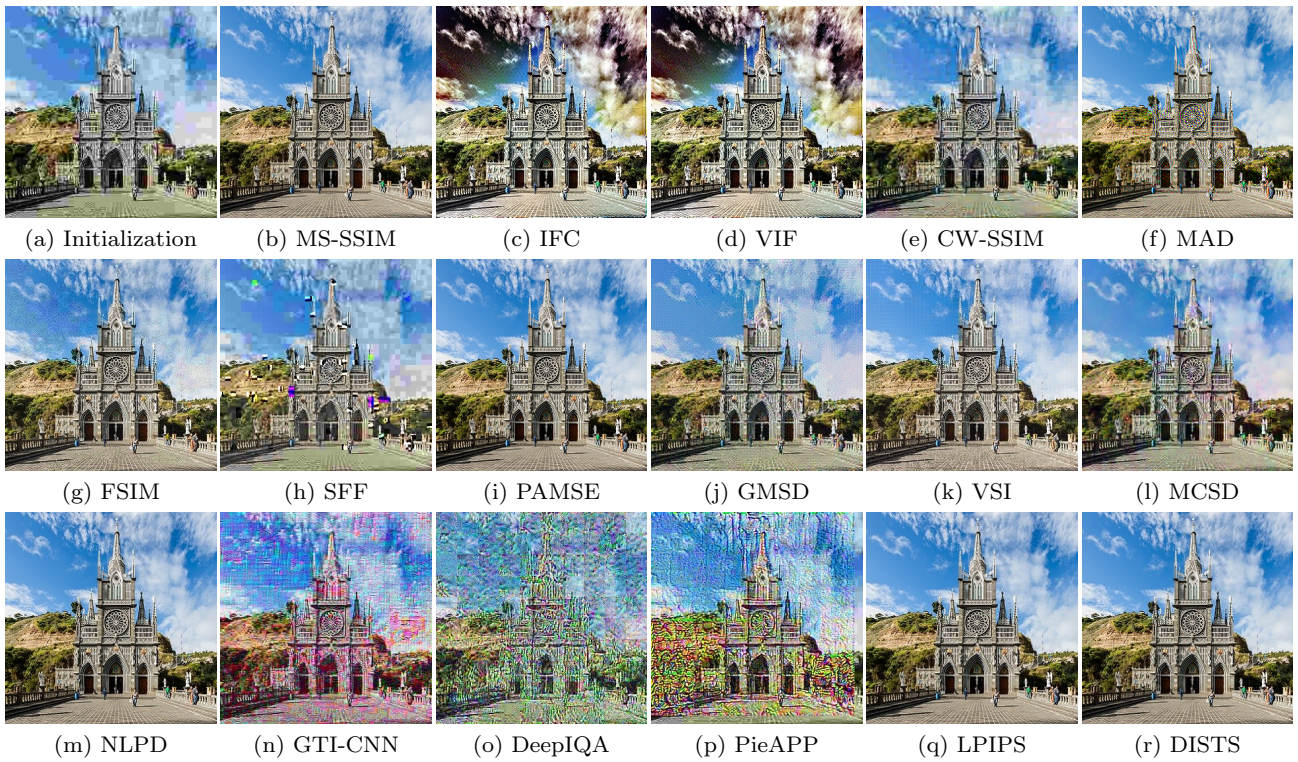


Fig. 2 Reference image recovery results, initialized from (a) a JPEG compressed version of the reference image. (b)-(r) Images recovered by optimizing different IQA models.

5. MAD (Larson and Chandler, 2010), the Most Apparent Distortion measure, explicitly models adaptive strategies of the HVS. Specifically, a detection-based strategy is employed for near-threshold distortions, and an appearance-based strategy is activated if the distortions are clearly visible.
6. FSIM (Zhang et al., 2011), the Feature SIMilarity index, computes quality estimates based on phase congruency (Kovesi, 1999) as the primary feature, and incorporates the gradient magnitude as the complementary feature. It also supplies a color version by making quality measurements from chromatic components.
7. GMSD (Xue et al., 2014), the Gradient Magnitude Similarity Deviation, computes pixel-wise gradient magnitude similarity followed by standard deviation (std) pooling. This pooling strategy is problematic: an image with large but constant local distortion yields an std of zero (indicating the best predicted quality).
8. VSI (Zhang et al., 2014), the Visual Saliency Induced quality index, assumes that the change of salient regions due to image degradation is closely related to the change of visual quality. It combines saliency magnitude similarity with gradient magnitude similarity, and demonstrates good quality prediction performance, especially for localized distortions, such as local patch substitution (Ponomarenko et al., 2015).
9. NLPD (Laparra et al., 2016), the Normalized Laplacian Pyramid Distance, mimics the nonlinear transformations of the early visual system: local luminance subtraction and local gain control, and combines these values using weighted ℓ_p -norms. The parameters are optimized to minimize the representation redundancies, instead of matching human judgments. NLPD has been successfully employed to optimize image rendering algorithms (Ma et al., 2015; Laparra et al., 2017), where the input reference image has a much higher dynamic range than that of the display. It has also been used to optimize a compression system (Ballé et al., 2016).
10. LPIPS (Zhang et al., 2018), the Learned Perceptual Image Patch Similarity model, computes the distance between deep representations of two images. The authors showed that feature maps of different DNN architectures have “reasonable” effectiveness in accounting for human perception of image quality. As LPIPS has many different configurations, we choose the default one based on the VGG network (Simonyan and Zisserman, 2015) with the weights learned from the BAPPS dataset (Zhang et al., 2018).
11. DISTS (Ding et al., 2020), the Deep Image Structure and Texture Similarity metric, is designed with explicit tolerance to texture resampling (e.g., replacing one patch of grass with another). DISTS is based on an injective mapping function built from a variant of the VGG network, and combines structure and texture similarity measurements between corresponding feature maps of the two images. It is sensitive to structural distortions, tolerant of texture resampling, and robust to mild geometric transformations.

We re-implemented all 11 of these models using PyTorch¹, and verified that our code could reproduce the published performance results for each model on the LIVE (Sheikh et al., 2006), CSIQ (Larson and Chandler, 2010), and TID2013 (Ponomarenko et al., 2015) databases (see Table 2 in Appendix A). We also modified grayscale-only models to accept color images, by computing scores on RGB channels separately and averaging them to obtain an overall quality estimate.

4 Perceptual Optimization of Standard Image Processing Tasks

We used each of the 11 full-reference IQA models as objective functions for optimizing the parameters of DNNs to solve four low-level vision tasks:

- image denoising,
- blind image deblurring,
- single image super-resolution,
- lossy image compression.

The parameters of each network are optimized to minimize an IQA measure over a database of corrupted and original image pairs via stochastic gradient descent. Implementations of all IQA models, as well as the DNNs for the four tasks, are available at <https://github.com/dingkeyan93/IQA-optimization>.

4.1 Image Denoising

Image denoising is a core application of classical image processing, and also plays an essential role in testing prior models of natural images. In its simplest form, one aims to recover an unknown clean image $x \in \mathbb{R}^N$ from an observed image y that has been corrupted by additive white Gaussian noise n of known variance σ^2 , i.e., $y = x + n$. Denoising algorithms can be roughly classified into spatial domain methods (e.g., Wiener filter (Weiner, 1950), bilateral filter (Tomasi and Manduchi, 1998) and collaborative filtering (Dabov et al.,

¹ <https://pytorch.org>

2007)), and wavelet transform methods (Donoho and Johnstone, 1995; Simoncelli and Adelson, 1996; Portilla et al., 2003). Later, sparsifying transforms (Elad and Aharon, 2006) and variants of nonlinear shrinkage functions have been directly learned from natural image data (Hel-Or and Shaked, 2008; Raphan and Simoncelli, 2008). In recent years, purely data-driven models based on DNNs (Zhang et al., 2017) have achieved new levels of performance.

Here, we constructed a simplified DNN, shown in Fig. 3, inspired by the EDSR network (Lim et al., 2017). The network was trained to estimate the noise (which is then subtracted from the observation to yield a denoised image), by minimizing a loss function defined by

$$\ell(\phi) = D(y - f_\phi(y), x), \quad (2)$$

where D is an IQA measure and $f_\phi : \mathbb{R}^N \mapsto \mathbb{R}^N$ is the mapping of the DNN, parameterized by vector ϕ .

4.2 Blind Image Deblurring

The goal of image deblurring is to restore a sharp image x from a blurry observation y , which can occur due to camera defocus or motion, and/or the motion of objects in a scene. The observation process is usually described by

$$y = Kx + n, \quad (3)$$

where $K \in \mathbb{R}^{N \times N}$ denotes a spatially-varying linear kernel. Blind deblurring refers to the problem in which the blur kernel is unknown. Most early methods, e.g., the classical Lucy-Richardson algorithm (Richardson, 1972; Lucy, 1974), focused on non-blind deblurring where the blur kernel is given. Successful blind deblurring methods, such as (Fergus et al., 2006; Pan et al., 2016), rely heavily on statistical priors of natural images and geometric priors of blur kernels. With the success of deep learning, many DNN-based approaches (Tao et al., 2018; Kupyn et al., 2018) attempt to directly learn the mapping function for blind deblurring without explicitly estimating the blur kernel. Here we also adopted this “kernel-free” approach to train a DNN for image deblurring in an end-to-end fashion. We employed the same network architecture used in denoising (see Fig. 3) with the same loss function (Eq. (2)).

4.3 Single Image Super-Resolution

Single image super-resolution aims to enhance the resolution and quality of a low-resolution image, which can

be modelled by

$$y = PKx + n, \quad (4)$$

where P denotes downsampling by a factor of β . This is an ill-posed problem, as downsampling is a projection onto a lower-dimensional subspace. Early attempts exploited sampling theory (Li and Orchard, 2001) or natural image statistics (Sun et al., 2008). Later methods focused on learning mapping functions between the low-resolution and high-resolution images through sparse coding (Yang et al., 2010), locally linear regression (Timofte et al., 2013), self-exemplars (Huang et al., 2015), etc. Since 2014, DNN-based methods have come to dominate this field as well (Dong et al., 2014). An efficient method of constructing a DNN-based mapping is to first extract features from the low-resolution input and then upscale them with sub-pixel convolution (Shi et al., 2016; Lim et al., 2017). Here, we followed this method in constructing a DNN-based function $f : \mathbb{R}^{\lfloor \frac{N}{\beta} \rfloor} \mapsto \mathbb{R}^N$, with architecture specified in Fig. 4. The loss is specified by

$$\ell(\phi) = D(f_\phi(y), x). \quad (5)$$

4.4 Lossy Image Compression

Data compression involves finding a more compact data representation from which the original image can be reconstructed. Compression can be either lossless or lossy. Here we followed a prevailing scheme in lossy image compression - transform coding, which consists of transformation, quantization, and entropy coding. Traditional image compression methods (e.g., the most widely used standard - JPEG) use a fixed linear transform for all bit rates. More recently, many researchers have demonstrated the visual benefits of nonlinear transforms, especially DNN-based learnable ones that are capable of adapting their parameters to different bitrate budgets. In this paper, we constructed two DNNs for analysis and synthesis transforms, respectively, as shown in Fig. 5. The analysis transform f_a maps the image to a latent feature vector z , whose values are then quantized to L levels with the centers being $\{c_1, \dots, c_L\}$, where $c_i \in \mathbb{R}$ for $i = 1, \dots, L$. This quantized representation $\bar{z} = Q(f_a(x))$, is fed to the synthesis transform f_s to reconstruct the compressed image: $y = f_s(\bar{z})$. The quantization step has zero gradients almost everywhere, which prevents gradient-based training via backpropagation (Ballé et al., 2017). Hence, we used a soft differentiable approximation (Mentzer et al., 2018)

$$\bar{z}_i = Q(z_i) = \sum_{j=1}^L \frac{\exp(-s(z_i - c_j)^2)}{\sum_{k=1}^L \exp(-s(z_i - c_k)^2)} c_j \quad (6)$$

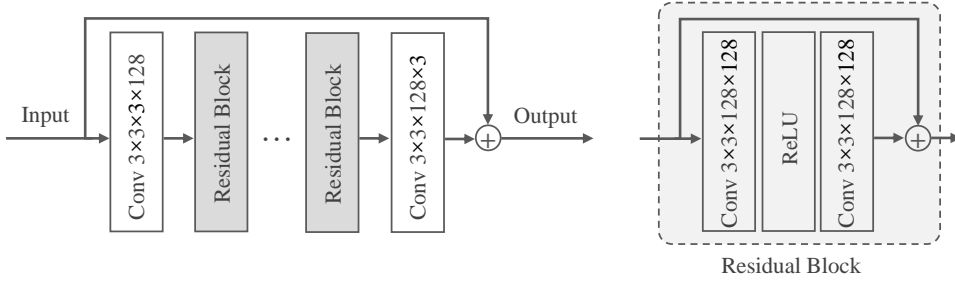


Fig. 3 Network architecture used for denoising and deblurring. Apart from initial and final convolutional blocks, it contains 16 residual blocks, each consisting of two convolutions and a halfwave rectifier (ReLU). Conv $h \times w \times c_{in} \times c_{out}$ indicates affine convolution with filter size $h \times w$, over c_{in} input channels, producing c_{out} output channels.




Fig. 4 Network architecture used for super-resolution, containing 16 residual blocks followed by two upsampling modules, each composed of an upsampler (factor of 2, using nearest-neighbor interpolation) and a convolution.




Fig. 5 Network architecture used for lossy image compression, which includes an analysis transformation f_a , a quantizer Q , and a synthesis transformation f_s . f_a is comprised of n blocks, each with a convolution and downsampling (stride) by 2 followed by two residual blocks. After the last block, another convolution layer with m filters is added to produce the internal code representation, the values of which are then quantized by Q . f_s consists of a cascade that is mirror-symmetric to f_a , with nearest-neighbor interpolation used to upsample the feature maps.

to backpropagate gradients during training, where the scale parameter s controls the approximation level of quantization.

In lossy image compression, the objective function is a weighted sum of two terms that quantify the coding cost and the reconstruction error, respectively:

$$\ell = \lambda H[\bar{z}] + \mathbb{E}[D(y, x)]. \quad (7)$$

The first term is typically the entropy (Shannon, 1948) of the discrete codes \bar{z} , which provides a lower bound on the bitrate for transmitting the quantized coefficients. The second term is the distortion between the compressed image y and the original image x , which we quantified with an IQA model D . The Lagrange multiplier λ controls the rate-distortion trade-off. Due to substantially different scales of IQA models, we would need

to manually adjust λ for each model in order to enable fair comparison at similar bitrates. To avoid this, following Agustsson et al. (2019), we set $\lambda = 0$ in Eq. (7), and controlled an upper bound on bitrate

$$H(\bar{z}) \leq \dim(\bar{z}) \log_2(L) \quad (8)$$

by adjusting the architecture of f_s (i.e., the dimension of \bar{z}) and the number of quantization levels L in Q . This elimination of the entropy from the objective also means that we did not need to continually re-estimate the probability mass function $P(\bar{z})$, which varies with changes in the network parameters. The optimization objective in Eq. (7) is reduced to

$$\ell(\phi, \psi) = \mathbb{E} \left[D \left(f_{s,\psi} \left(Q \left(f_{a,\phi}(x) \right) \right), x \right) \right], \quad (9)$$

where ϕ and ψ are the parameters of f_a and f_s , respectively. The expectation is approximated by averaging over mini-batches of training images.

5 Implementation Issues

In this section, we present in detail the training of our DNN-based computational models for the four low-level vision tasks, and the subjective testing to collect human rating of the optimized images as the ground truth.

5.1 Model Training

For denoising, we fixed the noise std to $\sigma = 50$. For deblurring, we simulated various kernels with different motion patterns and blur levels as in Kupyn et al. (2018). For super-resolution, we generated low-resolution images by downsampling high-resolution images by a factor of $\beta = 4$ using bicubic interpolation. For compression, we set the number of quantization levels to $L = 2$ with centers $\{-1, 1\}$, the quantization scale parameter to $s = 1$, the number of downsampling stages to $n = 4$, and the number of output channels of f_a to $m = 64$. This leads to a maximum of $\frac{H(\bar{z})}{W \times H} \leq \frac{W \times H}{2^4 \cdot 2^4} \cdot 64 \cdot \log_2(2) / (W \times H) = 0.25$ bpp (bits per pixel).

We chose the 4,744 high-quality images in the Waterloo Exploration Database (Ma et al., 2017b) as reference images. Training was performed in two stages. In the first stage, we pre-trained a network using MAE as the loss function for all four tasks. In the second stage, we fine-tuned the network parameters by optimizing the desired IQA model. Pre-training brings several advantages. First, a number of models are sensitive to initializations (e.g., CW-SSIM, MAD, FSIM, GMSD, and VSI) and pre-training yields more reasonable optimization results (also validated in the task of reference image recovery). Second, models that require backpropagating gradients through multiple stages of computation (e.g., LPIPS and DISTS) converge much faster. Third, it helps us to test whether the recently proposed IQA models lead to consistent perceptual gains on top of MAE, a special case of the simple ℓ_p -norm distance.

For each training stage of the four tasks, we used the Adam optimization package (Kingma and Ba, 2015) with a mini-batch size of 16 and an initial learning rate of 10^{-4} , which decays linearly by a factor of 2 for every 100K iterations, and we set the maximum number of iterations to 500K. We randomly extracted patches with the size of $192 \times 192 \times 3$ during training, and tested on 20 independent images from the DIV2K validation set (see Fig. 6). Training took roughly 1,000 GPU hours (measured using an NVIDIA GTX 2080 device) for a



Fig. 6 Test images (from the validation set of DIV2K) used in the subjective experiment.

total of $4 \times 11 = 44$ models. Special treatment (e.g., gradient clipping) was given to some IQA models (e.g., FSIM and VSI) to facilitate training and convergence.

Generally, it can be difficult to stabilize the training of DNNs to convergence, especially given that the gradients of different IQA models exhibit idiosyncratic behaviors. Fortunately, a simple criterion exists to test the validity of the optimization results: for a given low-level vision task, the DNN optimized for the IQA measure D_i should produce the best result (averaged over an independent set of images) in terms of D_i itself, when comparing to DNNs optimized for $\{D_j\}_{j \neq i}$. Fig. 7 shows ranking of results generated by networks optimized for each of the 11 IQA models (corresponding to one column in one subfigure) on the DIV2K validation set (Timofte et al., 2017), where 1 and 11 indicate the best and worst rankings, respectively. By inspecting the diagonal elements of the four matrices, we conclude that 43 out of 44 models satisfy the criterion, verifying the rationality of our training procedures. The only exception is when MAE is the optimization goal and NLPD (Laparra et al., 2016) is the evaluation measure for the deblurring task. Nevertheless, MAE ranks its own results the second place. As shown in Sec. 6.2, the

MAE	1	2	2	3	5	4	4	4	2	2	7
MS-SSIM	2	1	3	4	3	3	2	3	3	3	6
VIF	10	11	1	11	11	11	11	10	11	10	11
CW-SSIM	4	5	6	1	2	9	7	8	5	6	4
MAD	3	4	5	5	1	5	5	5	4	4	2
FSIM	5	6	7	8	4	1	3	2	6	8	5
GMSD	11	10	11	10	10	10	1	11	10	11	10
VSI	7	9	8	9	8	2	6	1	7	9	8
NLPD	6	3	4	2	9	8	8	9	1	7	9
LPIPS	9	7	9	6	6	6	9	7	8	1	3
DISTS	8	8	10	7	7	7	10	6	9	5	1
	MAE	MS-SSIM	VIF	CW-SSIM	MAD	FSIM	GMSD	VSI	NLPD	LPIPS	DISTS

(a) Denoising

MAE	2	3	9	8	9	9	10	10	2	8	6
MS-SSIM	4	1	8	7	7	8	7	8	4	5	5
VIF	11	11	1	11	11	11	11	11	11	10	10
CW-SSIM	3	2	6	1	5	7	8	7	3	4	4
MAD	7	4	4	2	1	5	4	5	7	3	3
FSIM	6	7	5	5	4	1	3	2	8	7	9
GMSD	10	10	11	10	8	6	1	4	10	11	11
VSI	5	9	7	6	6	2	2	1	5	6	8
NLPD	1	5	10	9	10	10	9	9	1	9	7
LPIPS	8	6	2	3	2	3	5	3	6	1	2
DISTS	9	8	3	4	3	4	6	6	9	2	1
	MAE	MS-SSIM	VIF	CW-SSIM	MAD	FSIM	GMSD	VSI	NLPD	LPIPS	DISTS

(b) Deblurring

MAE	1	2	3	4	5	5	3	8	2	6	7
MS-SSIM	3	1	4	5	4	3	2	4	3	3	6
VIF	11	11	1	11	11	10	11	10	11	11	11
CW-SSIM	8	9	11	1	10	11	10	11	5	10	10
MAD	4	4	5	2	1	7	6	6	4	4	3
FSIM	9	8	10	9	8	1	7	2	10	9	4
GMSD	10	10	9	10	9	9	1	5	9	8	9
VSI	6	7	8	8	7	2	5	1	8	7	5
NLPD	2	3	2	3	6	6	4	9	1	5	8
LPIPS	5	5	6	7	2	4	8	3	6	1	2
DISTS	7	6	7	6	3	8	9	7	7	2	1
	MAE	MS-SSIM	VIF	CW-SSIM	MAD	FSIM	GMSD	VSI	NLPD	LPIPS	DISTS

(c) Super-resolution

MAE	1	2	5	4	6	5	3	5	3	4	5
MS-SSIM	2	1	2	2	2	3	2	3	2	3	3
VIF	11	11	1	11	11	10	11	10	11	10	10
CW-SSIM	9	10	11	1	10	11	10	11	10	11	11
MAD	3	3	6	5	1	6	7	6	4	7	4
FSIM	4	5	4	8	5	1	5	2	6	5	7
GMSD	10	9	8	9	9	9	1	8	7	9	9
VSI	5	6	7	7	7	2	4	1	5	8	8
NLPD	8	4	3	3	8	8	6	9	1	6	6
LPIPS	7	8	9	10	4	4	9	4	9	1	2
DISTS	6	7	10	6	3	7	8	7	8	2	1
	MAE	MS-SSIM	VIF	CW-SSIM	MAD	FSIM	GMSD	VSI	NLPD	LPIPS	DISTS

(d) Compression

Fig. 7 Objective ranking of the final results in the four tasks. Vertical axis indicates IQA models used to train the networks, and horizontal axis indicates IQA models used to evaluate performance. The numbers of 1 to 11 indicate the rank order from the best to the worst.

resulting images from MAE and NLPD look visually similar.

5.2 Subjective Testing

We conducted a comprehensive subjective study to acquire human opinions of the perceptual optimization results. A two-alternative forced choice (2AFC) method was employed, allowing differentiation of fine-grained quality variations. Specifically, subjects were asked to choose which of two images has better perceived quality. The original image was also shown for reference. Subjects were allowed unlimited viewing time, and were free to adjust their viewing distance. Our customized graphical user interface also allows them to zoom in/out any portion of the two images for more careful visual comparison. We formed a total of $\binom{11}{2} \times 4 \times 20 = 4400$ paired

comparisons for 11 IQA models, 4 tasks, and 20 test images. To reduce fatigue, we performed the experiment in multiple sessions, each consisting of 500 randomly selected comparisons, and allowed subjects to take a break at anytime during the session. Subjects were encouraged, but not required, to participate in multiple sessions.

An outlier detection feature was also added: we included 5 pairs where one image was of unambiguously better quality (e.g., the original and noisy images). Our intention was to discard the voting results of subjects who failed in more than one of these pairs. We gathered data from 25 subjects with general background knowledge of image processing and computer vision, but were otherwise naïve to the purpose of this study. Voting results of all subjects turned out to be valid. In total, each image pair was evaluated by at least 5 subjects,

(a)	MS-SSIM MAE 0.70 0.65	MAD LPIPS 0.45 0.45	DISTS 0.39	NLPD 0.37	CW-SSIM 0.36	VSI VIF FSIM -0.44 -0.51 -0.58	GMSD -2.04
(b)	DISTS LPIPS 3.23 3.10	MAD MS-SSIM MAE 0.48 0.32 0.20	CW-SSIM 0.16	VIF NLPD -0.79 -0.94	FSIM VSI -1.54 -1.73	GMSD -2.75	
(c)	DISTS 2.50	LPIPS 1.88	MS-SSIM MAE 1.20 1.02	NLPD MAD 0.65 0.53	FSIM -0.70	VIF VSI GMSD -1.37 -1.81 -1.85	CW-SSIM -2.04
(d)	DISTS LPIPS 2.61 2.35	MS-SSIM MAE 1.58 1.53	MAD 0.68	NLPD 0.29	FSIM -0.37	VIF VSI GMSD -1.64 -2.00 -2.06	CW-SSIM -4.26

Fig. 8 Subjective ranking of the final results in the four tasks, based on human opinion scores. (a) Denoising. (b) Deblurring. (c) Super-resolution. (d) Compression. The optimization performance of IQA models is ranked in the descending order from left to right. Below each model is the global ranking score (larger is better). Models with the same colored box have statistically insignificant performance.

and each IQA model was ranked over 1,000 times for each vision task. The numbers of valid responses for denoising, deblurring, super-resolution, and compression were 6,516, 6,471, 6,473, and 6,540, respectively.

6 Experimental Results

In this section, based on subjective data, we conducted a quantitative comparison of the IQA models through the lens of perceptual optimization, yielding observations that are difficult to obtain from existing IQA databases. We also qualitatively compared the visual results associated with the IQA models. Last, we combined a top-performing IQA model with adversarial loss (Goodfellow et al., 2014) to test whether additional perceptual gains could be obtained in blind image deblurring.

6.1 Quantitative Results

We employed the BradleyTerry model (Bradley and Terry, 1952) to convert paired comparison results to global rankings. This probabilistic model assumes that the visual quality of the k -th test image optimized for the i -th IQA model, q_i^k , follows a Gumbel distribution with location μ_i^k and scale s . Assuming independence between q_i^k and q_j^k , the difference $q_i^k - q_j^k$ is a logistic random variable, and therefore $p_{ij}^k = P(q_i^k \geq q_j^k)$ can be computed using the logistic cumulative distribution function:

$$p_{ij}^k = P(q_i^k - q_j^k \geq 0) = \frac{\exp(\mu_i^k/s)}{\exp(\mu_i^k/s) + \exp(\mu_j^k/s)}, \quad (10)$$

where s is usually set to 1, leading to a simplified expression:

$$p_{ij}^k = \frac{e^{\mu_i^k}}{e^{\mu_i^k} + e^{\mu_j^k}}. \quad (11)$$

As such, we may obtain the negative log-likelihood of our pairwise count matrix W^k :

$$\ell(\mu^k | W^k) = \sum_{i=1}^M \sum_{\substack{j=1 \\ j \neq i}}^M \left(w_{ij}^k \log \left(e^{\mu_i^k} + e^{\mu_j^k} \right) - w_{ij}^k \mu_i^k \right), \quad (12)$$

where w_{ij}^k represents the number of times that D_i is preferred over D_j for the k -th test image. For each of the four low-level vision tasks, we minimized Eq. (12) iteratively using gradient descent to obtain the optimal estimate $\hat{\mu}^k$. We averaged $\hat{\mu}^k$ over the 20 test images, resulting in four global rankings of perceptual optimization performance, as shown in Fig. 8. It is clear that MS-SSIM (Wang et al., 2003) and MAE are superior to the other IQA models in the task of denoising, whereas DNN-based measures DISTS (Ding et al., 2020) and LPIPS (Zhang et al., 2018), outperform the others in the tasks of deblurring, super-resolution, and compression. Thus, there is no single IQA model that performs best across all tasks. We ascribe this to differences in the nature of the tasks: denoising requires distinguishing signal and noise, whereas deblurring, super-resolution and compression must recover missing details conditioned on the degraded information. Therefore, MS-SSIM and MAE that are known to prefer smooth appearances excel at denoising, despite close quantitative performance for most IQA models; DISTS and LPIPS that explicitly represent aspects of fine texture are superior for the remaining three tasks. Finally, it is important to note that many recent models, despite their impressive abilities to explain existing IQA databases, do not offer additional perceptual gains over MAE (and may even reduce perceptual quality).

To investigate whether the optimization results of the IQA models are statistically significant, we conducted an independent two-sample t -test. The null hypothesis is that the ranking scores $\{\mu_i^k\}_{k=1}^{20}$ for D_i and




Fig. 9 Denoising results on two regions cropped from an example image, using a DNN optimized for different IQA models.

$\{\mu_j^k\}_{k=1}^{20}$ for D_j come from the same normal distribution with unknown variance. When the test accepts the null hypothesis at the $\alpha = 5\%$ significance level, the two IQA models belong to the same group, and have the statistically indistinguishable performance. The grouping results are shown in Fig. 8, from which we find that the perceptual gains of MS-SSIM over MAE are statistically insignificant on all four tasks, which is quite surprising because MS-SSIM is widely regarded as a much better perceptual IQA model than MAE. Relying on similar sets of VGG features (Simonyan and Zisserman, 2015), DISTs and LPIPS also achieve similar performance, except for the super-resolution task where the former is statistically better.

By computing the SRCC between objective model rankings (in Fig. 7) and subjective human rankings (in Fig. 8), we are able to compare the algorithm-level performance of the 11 IQA models. We find from the Table 1 that there is a lack of correlation between model predictions and human judgments for the majority of IQA methods. DISTs and LPIPS tend to rank the images with complex model-dependent distortions in a more perceptually consistent way. We refer the interested readers to Appendix A for image-level comparison

Table 1 SRCC of objective ranking scores from the IQA models against subjective ranking scores

IQA Model	Denoising	Deblurring	Super-resolution	Compression
MAE	0.527	0.164	0.309	0.455
MS-SSIM	0.564	0.127	0.455	0.346
VIF	0.273	0.600	0.418	0.018
CW-SSIM	0.382	0.418	0.091	0.018
MAD	0.418	0.455	0.346	0.382
FSIM	0.236	0.054	0.091	0.127
GMSD	0.091	0.018	0.127	0.127
VSI	0.164	0.018	0.018	0.091
NLPD	0.491	0.127	0.200	0.309
LPIPS	0.709	0.855	0.782	0.782
DISTs	0.346	0.891	0.782	0.855

on several IQA databases dedicated to various low-level vision problems.

6.2 Qualitative Results

In this section, we show some visual examples produced by each IQA-optimized model, summarize the main types of distortions, and diagnose the shortcomings of the competing full-reference IQA methods.

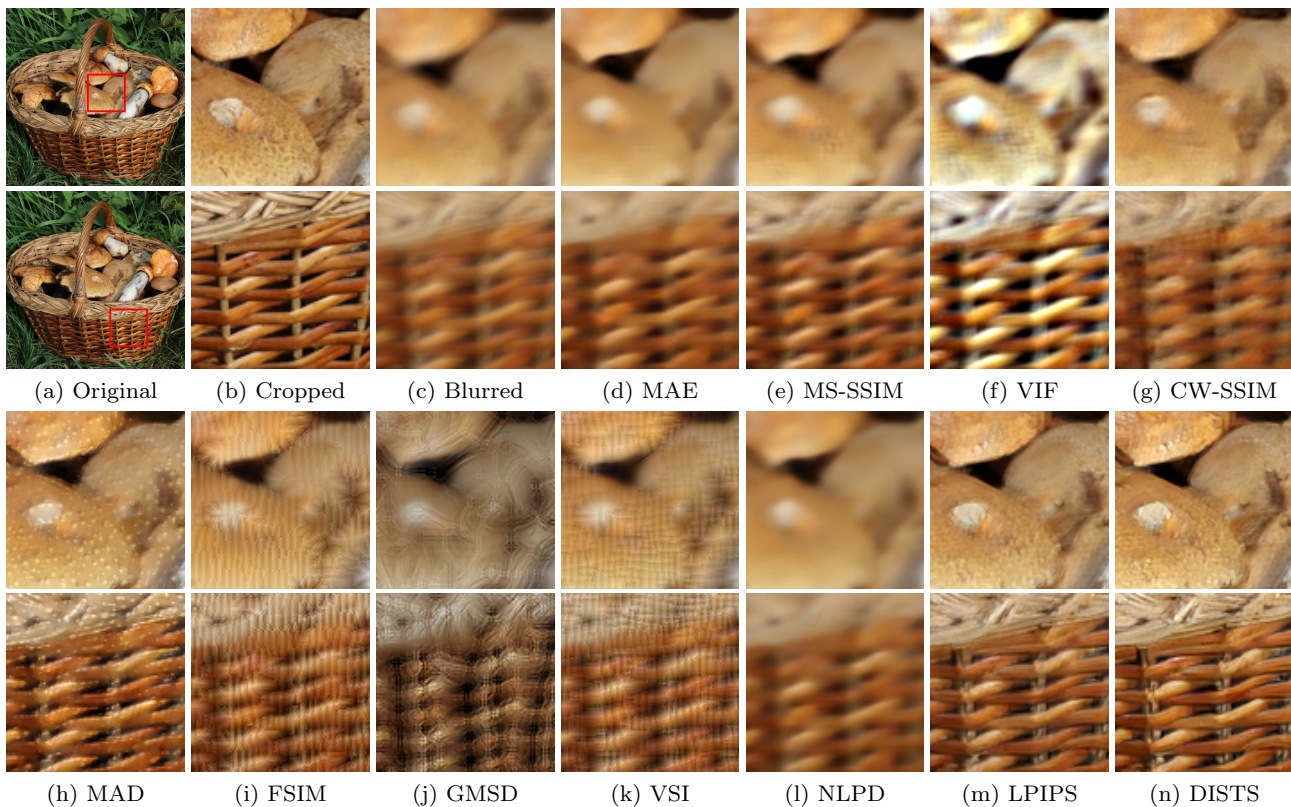


Fig. 10 Deblurring results for two regions cropped from an example image, using a DNN optimized for different IQA models.

Fig. 9 shows the *denoising* results of the “cat” image. It is not hard to observe that MAE, MS-SSIM, and NLPD do a good job in denoising flat regions, but tend to over-smooth texture regions. VIF encourages detail enhancement, leading to artificial local contrast, while GMSD produces a relatively dark appearance presumably because it discards local luminance information. Moreover, the results of FSIM and VSI exhibit noticeable artifacts. LPIPS and DISTIS preserve fine details, but may not fully remove noise in smooth regions, mistaking the remaining noise as visually plausible texture. Overall, traditional IQA models MAE and MS-SSIM denoise images with various content variations robustly, keeping high-frequency information loss within the acceptable range. This may explain why they are still the dominant objective functions in this task.

Fig. 10 shows the *deblurring* results of the “basket” image. We find that most IQA methods fail, but in different ways. Specifically, the results of MAE, MS-SSIM, CW-SSIM, and NLPD are still quite blurry. FSIM, GMSD, and VSI generate severe ringing artifacts. VIF again fails to adjust the local contrast. MAD exhibits undesirable dot artifacts, although the main structures are sharp. LPIPS succeeds in deblurring this example, while DISTIS produces a result that is closest to the

original. This is consistent with current state-of-the-art deblurring results, generated by incorporating comparison of the later stages of the VGG features into the loss.

Fig. 11 shows the *super-resolution* results of the “corner tower” image. Again, MAE, MS-SSIM, NLPD, and especially CW-SSIM produce somewhat blurry images, without recovering fine details that contain high-frequency information. MAD, FSIM, GMSD, and VSI are able to generate some “structures”, but these are perceived as unpleasant model-dependent artifacts. Benefiting from its texture synthesis capability, DISTIS has the potential to super-resolve perceptually plausible fine details, although they differ from those of the original image.

Fig. 12 shows the *compression* results of the “airplane” image at 0.24 ± 0.01 bpp. A JPEG image, compressed to 0.25 bpp, suffers from block and blur artifacts. Overall, the main structures of the original image are well preserved for most IQA models, but the fine details (e.g., the lawn) have to be discarded at this low bitrate, or are poorly synthesized as other forms of distortions. VIF reconstitutes a dreary image with over-enhanced global contrast, and CW-SSIM superimposes wavy artifacts on the underlying image. Dot and ring-




Fig. 11 Super-resolution results for two cropped regions from an example image, using a DNN optimized for different IQA models.

ing artifacts are again apparent in the results of MAD and VSI, respectively. The image by NLPD is blurred and red shifted. LPIPS and DISTIS have more potentials in synthesizing textures that are visually similar to the original, outperforming other IQA models in this task.

Now, we summarize and diagnose the “novel” artifacts created during perceptual optimization, which are not typically seen in traditional image databases for the purpose of quality assessment.

- *Blurring* is a frequently seen distortion type in all of the four vision tasks, and is mainly caused by error visibility methods (e.g., MAE and NLPD) and structural similarity methods (e.g., MS-SSIM), which rely on simple injective mappings. Specifically, MAE and SSIM work directly with pixels, and NLPD transforms the input image to a multiscale overcomplete representation using a single stage of local mean subtraction and divisive normalization. Under strict constraints imposed by the vision tasks, they prefer to make a more conservative estimate, producing something akin to a superposition of all possible outcomes with sharp structures, as would occur with a more conventional loss - MSE.
- *Ringing* is a high-frequency distortion type that often occurs in the optimized images by FSIM, VSI and GMSD (see Fig. 10 (i) - (k)). One common characteristic of the three models is that they rely heavily (in some cases, solely) on local gradient magnitude for feature similarity comparison, underweighting (or abandoning) other perceptually important features (such as local luminance and local phase). This creates enormous “shortcuts” that DNN-based computational models can take to generate distortions with similar local gradient statistics.
- *Dot* patterns are typical in the optimization results of MAD, which extracts lower-order image statistics from responses of Gabor filters at multiple scales and orientations. The resulting set of statistical measurements is insufficient to summarize natural image structures that exhibit higher-order dependencies. Therefore, MAD is “blind” to distortions that satisfy the same set of statistical constraints, and gives the optimized distorted image a high-quality score.
- *Over-enhancement* of local image contrast is encouraged by VIF, which, in most of our experiments, causes significant quality degradation. We believe this arises because VIF does not fully respect reference information when normalizing the covariance




Fig. 12 Compression results for two cropped regions from an example image, using a DNN optimized for different IQA models.

term. Specifically, only the second-order statistics of the reference image are used to construct the normalization factor. By incorporating the same statistics computed from the distorted image into normalization, the problem of over-enhancement may be alleviated. In general, quality assessment of image enhancement is a challenging problem (Fang et al., 2015; Wang et al., 2015), and to the best of our knowledge, all existing full-reference IQA models fail to reward properly-enhanced cases, while penalizing over-enhanced cases.

- *Luminance and color* artifacts are perceived in final images that are associated with many IQA models. Two causes seem plausible. First, methods such as GMSD discard luminance information, leaving a huge “null space” to accommodate luminance distortions. Second, methods such as MS-SSIM and NLPD are originally designed for grayscale images only. Applying them to RGB channels separately fails to take into account saturation (color contrast). Transforming to a perceptually better color space, and making use of knowledge of color distortions (Rajashekar et al., 2009) offers an opportunity for improvement.

6.3 Combining with Adversarial Loss

In the field of image restoration and generation, many state-of-the-art algorithms are based on adversarial training (Goodfellow et al., 2014), demonstrating impressive capability of synthesizing realistic visual content. The output of the adversarial loss is the probability of an image being computer-generated, and therefore is of low relevance to perceived quality of that image. In other words, the adversarial loss is a poor no-reference IQA model, at the level of images (verified by our experiments on the LIVE database). However, it may be a “good” one in the algorithm level, meaning that given a set of images generated by a computational method and another set of natural photographic images, the average probability on the combined set quantitatively measures the capability of the method in generating realistic high-quality images. In this subsection, we explored the combination of the adversarial loss and top-performing IQA measures for additional perceptual gains.

We chose the task of blind image deblurring, and fine-tuned a state-of-the-art model - DeblurGAN-v2 (Kupyn et al., 2019). According to our perceptual optimization results, we selected the best-performing IQA model




Fig. 13 Deblurring examples obtained by the original DeblurGAN-v2 (under the configuration of Inception-ResNet) and the fine-tuned using the DISTs IQA model.

(DISTs) for this experiment. We followed the same training strategy, and only replaced the loss function of the generator from

$$\ell_o = 0.5 \times \ell_{\text{MSE}} + 0.006 \times \ell_{\text{VGG}} + 0.01 \times \ell_{\text{Adv}} \quad (13)$$

to

$$\ell_n = \ell_{\text{DISTs}} + 0.001 \times \ell_{\text{Adv}}. \quad (14)$$

The first and second terms in Eq. (13) compute the MSE on pixels and responses of conv3_3 of VGG19 (Simonyan and Zisserman, 2015), respectively. ℓ_{Adv} denotes a variant of the adversarial loss (Kupyn et al., 2019). An immediate advantage of this replacement is that the number of hyperparameters is reduced, making manual hyperparameter adjustment easier. After fine-tuning, the average DISTs value decreases from 0.22 to 0.18 on the Köhler test dataset (Köhler et al., 2012). Fig. 13 shows two visual examples, from which we find

that the fine-tuned results have sharper edges and enhanced contrast, indicating that perceptual gains may be obtained by DISTs on the two examples.

7 Conclusions

We have conducted a comprehensive study of perceptual optimization of four low-level vision tasks, guided by 11 full-reference IQA models. This provides an alternative means of testing the perceptual relevance of IQA models in a more realistic setting, which we believe is an important complement to the conventional methodology for IQA model evaluation. Extensive subjective testing led to several useful findings. First, through perceptual optimization, we generated a number of novel distortions (different from those in existing IQA databases), which may easily fool many competing models. It should be noted that the emergence of specific distortions is in

principle dependent on the experimental setting (e.g., initialization strategy, model architecture, and optimization technique). Second, standard full-reference IQA models, MS-SSIM and MAE, will continue to play a central role in optimizing image processing systems due to their robustness and simplicity. Third, more recent IQA models with surjective mappings may still be used to monitor image quality and to optimize the parameter settings of image processing methods, but in a limited and well-controlled space. Last, the two DNN-based models, LPIPS and DISTS seem to stand out in our experiments, but their high computational complexity and lack of interpretability may hinder their application.

Our work has interesting connections to two separate lines of research. First, inspired by the philosophy of “analysis by synthesis” (Grenander, 1970), Wang and Simoncelli (2008) introduced the maximum differentiation competition methodology to automatically synthesize images for efficiently comparing IQA models. However, the generated images may be highly unnatural, and therefore are of less practical importance. Ma et al. (2020) alleviated this issue by manually constraining the search space to a finite image set of current interest. Our approach can be seen as a mixture of the two methods, in the sense that the test images are automatically generated, and are well-controlled by specifying realistic vision tasks. Second, the existence of type II adversarial examples (Szegedy et al., 2013) has exposed the vulnerability of many computer vision algorithms, where a tiny change to the input that is imperceptible to the human eye would cause the algorithm to make classification mistakes. In our case, weaknesses in an IQA model are exposed through optimized images that serve as type I adversarial examples of the model: a significant change is made to the original image that substantially degrades its perceptual quality, but the model claims this image is of high quality.

Fusion of IQA models, which has not been investigated in this paper, may be a feasible way of developing a better IQA method in the context of perceptual optimization, as constantly practiced by researchers in related fields. The main technical difficulty is the lack of a principled way to combine IQA models of different perceptual scales.

Besides model fusion, a more promising design strategy for new IQA methods is to enforce a set of desirable properties. First, the transformation used in the IQA model should be perceptual, mapping the input images into a space where Euclidean distances match human measurements of image quality. This is in the same spirit that color scientists pursue perceptually uniform color spaces. Zhang et al. (2018) and Ding et al. (2020)

demonstrated that a cascade of linear convolution, down-sampling, and rectified nonlinearity optimized for high-level vision tasks may be a good candidate. Second, the IQA model should enjoy unique optima (i.e., the underlying mapping should be injective) to guarantee that images close to optimum are visually similar to the original. This criterion was respected by early models (e.g., MS-SSIM), but was largely overlooked in recent IQA model development. Third, the IQA model should be continuous and differentiable, with well-behaved gradients, to aid optimization in complex situations (e.g., training DNNs with millions of parameters). Last but not least, the IQA model should be computationally efficient, enabling real-time quality assessment and perceptual optimization.

Appendix A Perceptual Correlation Comparison of IQA Models

A conventional method for evaluating IQA models is to compute their agreement with subjective scores in one or more standardized IQA databases (e.g., LIVE (Sheikh et al., 2006), CSIQ (Larson and Chandler, 2010) or TID2013 (Ponomarenko et al., 2015)), consisting of artificially distorted images. Many existing IQA models achieved impressive correlation numbers on these databases (see Table 2), but their performance in assessing the perceptual quality of images produced by low-level vision algorithms has not been tested. In this appendix, we tested them on multiple human-rated image generation/restoration quality assessment databases, including

- A denoising database - FLT (Egiazarian et al., 2018), consisting of 300 filtered images from 75 grayscale texture images by the BM3D algorithm with different levels of noise suppression.
- Two motion deblurring databases - Liu13 (Liu et al., 2013) and Lai16 (Lai et al., 2016). Liu13 contains 240 synthetically blurred examples, each of which is deblurred by five algorithms. Lai16 synthesizes 100 non-uniformly and 100 uniformly blurred images, which are further deblurred by 13 algorithms.
- Two super-resolution databases - Ma17 (Ma et al., 2017a) and QADS (Zhou et al., 2019). The former has 1,620 super-resolved images from 9 methods, while the latter contains 980 images created by 21 methods.
- A dehazing database - SHRQ (Min et al., 2019), composed of 45 regular hazy images and 30 aerial hazy images, which are dehazed by 8 algorithms.

Table 2 Verification of results obtained by our PyTorch implementations of the involved IQA models, on three IQA databases. Numbers indicate SRCC values (reported in original publication / produced by our implementation). Bold indicates that these methods are computed on the grayscale images in their original versions; here we extended them to evaluate the RGB images by averaging the values across all channels

IQA Model	Grayscale			Color		
	LIVE	CSIQ	TID2013	LIVE	CSIQ	TID2013
MS-SSIM	0.951 / 0.951	0.906 / 0.886	0.786 / 0.782	0.931 / 0.932	0.902 / 0.886	0.801 / 0.816
CW-SSIM	0.786 / 0.781	0.745 / 0.738	0.673 / 0.680	0.741 / 0.747	0.744 / 0.744	0.709 / 0.725
VIF	0.964 / 0.963	0.911 / 0.911	0.677 / 0.676	0.957 / 0.957	0.894 / 0.894	0.654 / 0.654
NLPD	0.937 / 0.938	0.932 / 0.937	0.800 / 0.800	0.917 / 0.914	0.913 / 0.913	0.812 / 0.808
GMSD	0.960 / 0.960	0.950 / 0.950	0.804 / 0.804	0.949 / 0.948	0.937 / 0.934	0.830 / 0.823
MAD	0.967 / 0.960	0.947 / 0.941	0.781 / 0.773	0.954 / 0.951	0.937 / 0.935	0.758 / 0.740
FSIM	0.963 / 0.963	0.924 / 0.916	0.802 / 0.802	0.965 / 0.965	0.931 / 0.923	0.851 / 0.851
VSI	0.953 / 0.950	0.930 / 0.923	0.805 / 0.793	0.952 / 0.956	0.942 / 0.937	0.897 / 0.889
LPIPS	0.932 / 0.932	0.837 / 0.837	0.616 / 0.616	0.932 / 0.932	0.876 / 0.876	0.670 / 0.670
DISTS	0.942 / 0.942	0.905 / 0.905	0.764 / 0.764	0.954 / 0.954	0.929 / 0.929	0.830 / 0.830

- A depth image-based rendering database - Tian19 (Tian et al., 2018), consisting of 140 rendered images from 10 sequences with 7 methods.
- Two texture synthesis databases - SynTex (Golestaneh et al., 2015) and TQD (Ding et al., 2020). SynTex contains 105 synthesized textures using 5 algorithms from 21 reference textures. TQD has 10 reference textures, each with 15 variations, including 7 artificially distorted images, 4 example-based synthesized textures and 4 cropped subimages.
- A patch database - BAPPS (Zhang et al., 2018), which includes 26.9K image patches generated by colorization, video deblurring, frame interpolation, and super-resolution algorithms, respectively.

Table 3 shows the performance comparisons of 13 IQA methods in terms of SRCC. Considering the potential noisy judgements, BAPPS is evaluated using the 2AFC score: $pq + (1 - p)(1 - q)$, where p is the percentage of human votes and $q = \{0, 1\}$ is the vote of an IQA model. When q is closer to the majority of human votes, the 2AFC score is larger, indicating better performance (see Table 4). We find that the overall performance of all models is lower compared to that in the standard IQA databases (see Table 2), indicating the difficulty when applying to unseen distortions. Moreover, DNN-based measures are relatively better than knowledge-driven models in these application-oriented databases, but there is still significant room for improvement to meet the real-world challenges.

Fig. 14 and Fig. 15 show two quality assessment examples of compression and super-resolution, respectively. Here we only compared the most widely used measures: PSNR and SSIM, and the two that performed best both on optimization and assessment: LPIPS and DISTS. It is not surprising to find that PSNR and SSIM have poor correlations with human opinions, as they focus more on signal fidelity rather than percep-

tual quality. LPIPS and DISTS perform better, but the former is somewhat oversensitive to texture substitution (see Fig. 15). As many recent image restoration algorithms succeed in generating richer textures, DISTS holds much promise for use in quality assessment for such applications.

Acknowledgements The authors would like to thank all participants who contributed to our subjective study in the special period of coronavirus disease (COVID-19) outbreak.

References

- Agustsson E, Tschannen M, Mentzer F, Timofte R, Gool LV (2019) Generative adversarial networks for extreme learned image compression. In: IEEE Conference on Computer Vision and Pattern Recognition, pp 221–231
- Ballé J, Laparra V, Simoncelli EP (2016) End-to-end optimization of nonlinear transform codes for perceptual quality. In: Picture Coding Symposium, pp 1–5
- Ballé J, Laparra V, Simoncelli EP (2017) End-to-end optimized image compression. In: International Conference on Learning Representations, pp 1–27
- Ballé J, Minnen D, Singh S, Hwang SJ, Johnston N (2018) Variational image compression with a scale hyperprior. In: International Conference on Learning Representations, pp 1–47
- Bosse S, Maniry D, Müller KR, Wiegand T, Samek W (2018) Deep neural networks for no-reference and full-reference image quality assessment. *IEEE Transactions on Image Processing* 27(1):206–219
- Bradley RA, Terry ME (1952) Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika* 39(3/4):324–345

Table 3 SRCC comparison of IQA models on existing image generation/restoration quality assessment databases

IQA Model	Denoising	Deblurring		Super-resolution		Dehazing	Rendering	Texture synthesis	
	FLT	Lai16	Liu13	Ma17	QADS	SHRQ	Tian19	SynTEX	TQD
PSNR	0.183	0.301	0.803	0.592	0.360	0.740	0.536	0.114	0.233
SSIM	0.355	0.298	0.777	0.624	0.529	0.692	0.230	0.620	0.307
MS-SSIM	0.246	0.320	0.898	0.795	0.717	0.687	0.396	0.469	0.288
VIF	0.169	0.261	0.864	0.831	0.815	0.667	0.259	0.448	0.305
CW-SSIM	0.101	0.600	0.742	0.706	0.474	0.698	0.522	0.496	0.325
MAD	0.182	0.446	0.897	0.864	0.723	0.605	0.622	0.134	0.302
FSIM	0.555	0.297	0.921	0.747	0.687	0.695	0.476	0.093	0.176
GMSD	0.389	0.174	0.918	0.851	0.765	0.663	0.479	0.006	0.256
VSI	0.528	0.295	0.920	0.710	0.584	0.696	0.531	0.123	0.179
NLPD	0.151	0.323	0.853	0.732	0.591	0.608	0.463	0.483	0.271
PieAPP	0.629	0.601	0.786	0.771	0.849	0.725	0.298	0.709	0.713
LPIPS	0.457	0.347	0.867	0.788	0.669	0.777	0.311	0.663	0.392
DISTS	0.636	0.754	0.941	0.878	0.809	0.789	0.671	0.923	0.910

Table 4 2AFC score comparison of IQA models on the BAPPS dataset and the proposed dataset

IQA Model	BAPPS				Proposed			
	Color-ization	Video deblurring	Frame interpolation	Super-resolution	Denoising	Deblurring	Super-resolution	Compression
PSNR	0.624	0.590	0.543	0.642	0.627	0.518	0.612	0.689
SSIM	0.522	0.583	0.548	0.613	0.636	0.575	0.599	0.649
MS-SSIM	0.522	0.589	0.572	0.638	0.623	0.568	0.655	0.665
VIF	0.515	0.594	0.597	0.651	0.589	0.607	0.655	0.540
CW-SSIM	0.512	0.601	0.604	0.665	0.623	0.651	0.584	0.496
MAD	0.490	0.593	0.581	0.655	0.624	0.671	0.681	0.651
FSIM	0.573	0.590	0.581	0.660	0.522	0.490	0.525	0.563
GMSD	0.517	0.594	0.575	0.676	0.417	0.454	0.469	0.567
VSI	0.597	0.591	0.568	0.668	0.518	0.470	0.487	0.576
NLPD	0.528	0.584	0.552	0.655	0.622	0.514	0.629	0.652
PieAPP	0.594	0.582	0.598	0.685	0.625	0.734	0.744	0.822
LPIPS	0.625	0.605	0.630	0.705	0.657	0.788	0.768	0.834
DISTS	0.627	0.600	0.625	0.710	0.602	0.790	0.704	0.833

**Fig. 14** A visual quality assessment example of compression. (a) Uncompressed image. (b)-(h) are the compressed results optimized for indicated IQA models. The results of MAE and MS-SSIM outperform the others in terms of PSNR and SSIM, but are not among the best judged by humans. LPIPS and DISTS are more consistent with human perception in this example.

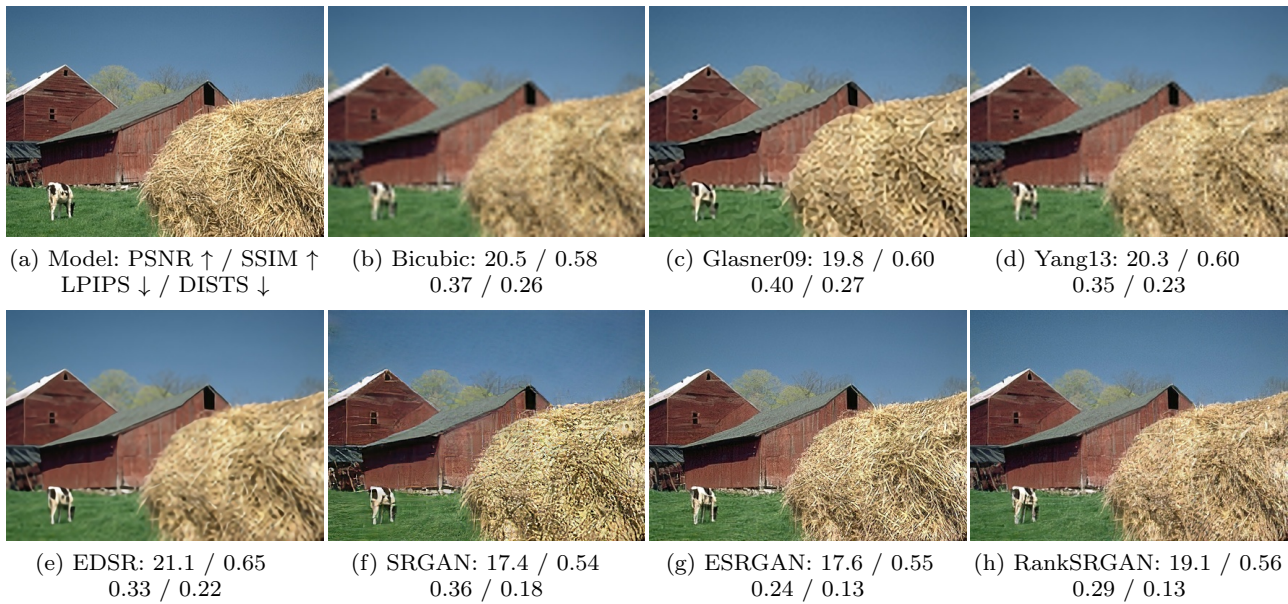


Fig. 15 A visual quality assessment example of super-resolution. (a) High-resolution image. (b)-(h) are the super-resolution results computed using bicubic interpolation, Glasner09 (Glasner et al., 2009), Yang13 (Yang and Yang, 2013), EDSR (Lim et al., 2017), SRGAN (Ledig et al., 2017), ESRGAN (Wang et al., 2018), and RankSRGAN (Zhang et al., 2019a), respectively. One can see that the GAN-based results (f)-(h) are visually superior to the others, contrary to the predictions of PSNR and SSIM. LPIPS indicates that the result (f) is worse than (d) and (e), in disagreement with visual inspection. DISTS is correlated well with human perception in this example.

Chang HW, Yang H, Gan Y, Wang MH (2013) Sparse feature fidelity for perceptual image quality assessment. *IEEE Transactions on Image Processing* 22(10):4007–4018

Channappayya SS, Bovik AC, Caramanis C, Heath RW (2008) SSIM-optimal linear image restoration. In: *IEEE International Conference on Acoustics, Speech, and Signal Processing*, pp 765–768

Dabov K, Foi A, Katkovnik V, Egiazarian K (2007) Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Transactions on Image Processing* 16(8):2080–2095

Daly SJ (1992) Visible differences predictor: An algorithm for the assessment of image fidelity. In: *Human Vision, Visual Processing, and Digital Display III*, pp 2–15

Ding K, Ma K, Wang S, Simoncelli EP (2020) Image quality assessment: Unifying structure and texture similarity. *CoRR* abs/2004.07728, [Online]. Available: <https://arxiv.org/abs/2004.07728>

Dong C, Loy CC, He K, Tang X (2014) Learning a deep convolutional network for image super-resolution. In: *European Conference on Computer Vision*, pp 184–199

Donoho DL, Johnstone IM (1995) Adapting to unknown smoothness via wavelet shrinkage. *Journal of the American Statistical Association* 90(432):1200–

1224

Egiazarian K, Ponomarenko M, Lukin V, Ieremeiev O (2018) Statistical evaluation of visual quality metrics for image denoising. In: *IEEE International Conference on Acoustics, Speech and Signal Processing*, pp 6752–6756

Elad M, Aharon M (2006) Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Transactions on Image Processing* 15(12):3736–3745

Fang Y, Ma K, Wang Z, Lin W, Fang Z, Zhai G (2015) No-reference quality assessment of contrast-distorted images based on natural scene statistics. *IEEE Signal Processing Letters* 22(7):838–842

Fergus R, Singh B, Hertzmann A, Roweis ST, Freeman WT (2006) Removing camera shake from a single photograph. *ACM Transactions on Graphics* 25(3):787–794

Girod B (1993) *What’s wrong with mean-squared error?* MIT Press, Cambridge, MA, USA

Glasner D, Bagon S, Irani M (2009) Super-resolution from a single image. In: *IEEE International Conference on Computer Vision*, pp 349–356

Golestaneh SA, Subedar MM, Karam LJ (2015) The effect of texture granularity on texture synthesis quality. In: *Applications of Digital Image Processing XXXVIII*, vol 9599, pp 356 – 361

- Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y (2014) Generative adversarial nets. In: Conference on Neural Information Processing Systems, pp 2672–2680
- Grenander U (1970) A unified approach to pattern analysis. In: *Advances in Computers*, vol 10, pp 175–216
- Hel-Or Y, Shaked D (2008) A discriminative approach for wavelet denoising. *IEEE Transactions on Image Processing* 17(4):443–457
- Huang JB, Singh A, Ahuja N (2015) Single image super-resolution from transformed self-exemplars. In: *IEEE Conference on Computer Vision and Pattern Recognition*, pp 5197–5206
- Johnson J, Alahi A, Li FF (2016) Perceptual losses for real-time style transfer and super-resolution. In: *European Conference on Computer Vision*, pp 694–711
- Jordan C (1881) Sur la série de fourier. *Comptes Rendus de l'Académie des Sciences Paris* 2:228–230
- Kingma DP, Ba J (2015) Adam: A method for stochastic optimization. In: *International Conference on Learning Representations*, pp 1–15
- Köhler R, Hirsch M, Mohler B, Schölkopf B, Harmeling S (2012) Recording and playback of camera shake: Benchmarking blind deconvolution with a real-world database. In: *European Conference on Computer Vision*, pp 27–40
- Kovesi P (1999) Image features from phase congruency. *Journal of Computer Vision Research* 1(3):1–26
- Kupyn O, Budzan V, Mykhailych M, Mishkin D, Matas J (2018) DeblurGAN: Blind motion deblurring using conditional adversarial networks. In: *IEEE Conference on Computer Vision and Pattern Recognition*, pp 8183–8192
- Kupyn O, Martyniuk T, Wu J, Wang Z (2019) DeblurGAN-v2: Deblurring (orders-of-magnitude) faster and better. In: *IEEE International Conference on Computer Vision*, pp 8878–8887
- Lai WS, Huang JB, Hu Z, Ahuja N, Yang MH (2016) A comparative study for single image blind deblurring. In: *IEEE Conference on Computer Vision and Pattern Recognition*, pp 1701–1709
- Laparra V, Ballé J, Berardino A, Simoncelli EP (2016) Perceptual image quality assessment using a normalized Laplacian pyramid. *Electronic Imaging* 2016(16):1–6
- Laparra V, Berardino A, Ballé J, Simoncelli EP (2017) Perceptually optimized image rendering. *Journal of the Optical Society of America A* 34(9):1511–1525
- Larson EC, Chandler DM (2010) Most apparent distortion: Full-reference image quality assessment and the role of strategy. *Journal of Electronic Imaging* 19(1):1–21
- Ledig C, Theis L, Huszár F, Caballero J, Cunningham A, Acosta A, Aitken A, Tejani A, Totz J, Wang Z, et al. (2017) Photo-realistic single image super-resolution using a generative adversarial network. In: *IEEE Conference on Computer Vision and Pattern Recognition*, pp 4681–4690
- Li X, Orchard MT (2001) New edge-directed interpolation. *IEEE Transactions on Image Processing* 10(10):1521–1527
- Lim B, Son S, Kim H, Nah S, Mu Lee K (2017) Enhanced deep residual networks for single image super-resolution. In: *IEEE Conference on Computer Vision and Pattern Recognition Workshop*, pp 136–144
- Liu A, Lin W, Narwaria M (2012a) Image quality assessment based on gradient similarity. *IEEE Transactions on Image Processing* 21(4):1500–1512
- Liu TJ, Lin W, Kuo CCJ (2012b) Image quality assessment using multi-method fusion. *IEEE Transactions on Image Processing* 22(5):1793–1807
- Liu Y, Wang J, Cho S, Finkelstein A, Rusinkiewicz S (2013) A no-reference metric for evaluating the quality of motion deblurring. *ACM Transactions on Graphics* 32(6)
- Lubin J (1993) *The use of psychophysical data and models in the analysis of display system performance*. MIT Press, Cambridge, MA, USA
- Lucy LB (1974) An iterative technique for the rectification of observed distributions. *Astronomical Journal* 79:745
- Ma C, Yang CY, Yang X, Yang MH (2017a) Learning a no-reference quality metric for single-image super-resolution. *Computer Vision and Image Understanding* 158:1–16
- Ma K, Yeganeh H, Zeng K, Wang Z (2015) High dynamic range image compression by optimizing tone mapped image quality index. *IEEE Transactions on Image Processing* 24(10):3086–3097
- Ma K, Duanmu Z, Wu Q, Wang Z, Yong H, Li H, Zhang L (2017b) Waterloo Exploration Database: New challenges for image quality assessment models. *IEEE Transactions on Image Processing* 26(2):1004–1016
- Ma K, Duanmu Z, Wang Z (2018) Geometric transformation invariant image quality assessment using convolutional neural networks. In: *IEEE International Conference on Acoustics, Speech, and Signal Processing*, pp 6732–6736
- Ma K, Liu X, Fang Y, Simoncelli EP (2019) Blind image quality assessment by learning from multiple annotators. In: *IEEE International Conference on Image Processing*, pp 2344–2348
- Ma K, Duanmu Z, Wang Z, Wu Q, Liu W, Yong H, Li H, Zhang L (2020) Group maximum differentiation competition: Model comparison with few samples. *IEEE*

- Transactions on Pattern Analysis and Machine Intelligence 42(4):851–864
- Mentzer F, Agustsson E, Tschannen M, Timofte R, Van Gool L (2018) Conditional probability models for deep image compression. In: IEEE Conference on Computer Vision and Pattern Recognition, pp 4394–4402
- Min X, Zhai G, Gu K, Zhu Y, Zhou J, Guo G, Yang X, Guan X, Zhang W (2019) Quality evaluation of image dehazing methods using synthetic hazy images. IEEE Transactions on Multimedia 21(9):2319–2333
- Pan J, Sun D, Pfister H, Yang MH (2016) Blind image deblurring using dark channel prior. In: IEEE Conference on Computer Vision and Pattern Recognition, pp 1628–1636
- Ponomarenko N, Jin L, Ieremeiev O, Lukin V, Egiazarian K, Astola J, Vozel B, Chehdi K, Carli M, Battisti F, Kuo CCJ (2015) Image database TID2013: Peculiarities, results and perspectives. Signal Processing Image Communication 30:57–77
- Portilla J, Strela V, Wainwright MJ, Simoncelli EP (2003) Image denoising using scale mixtures of Gaussians in the wavelet domain. IEEE Transactions on Image Processing 12(11):1338–1351
- Prashnani E, Cai H, Mostofi Y, Sen P (2018) PieAPP: Perceptual image-error assessment through pairwise preference. In: IEEE Conference on Computer Vision and Pattern Recognition, pp 1808–1817
- Rajashekar U, Wang Z, Simoncelli EP (2009) Quantifying color image distortions based on adaptive spatiochromatic signal decompositions. In: IEEE International Conference on Image Processing, IEEE, pp 2213–2216
- Raphan M, Simoncelli EP (2008) Optimal denoising in redundant representations. IEEE Transactions on Image Processing 17(8):1342–1352
- Richardson WH (1972) Bayesian-based iterative method of image restoration. Journal of Optical Society of America 62(1):55–59
- Safranek RJ, Johnston JD (1989) A perceptually tuned sub-band image coder with image dependent quantization and post-quantization data compression. In: IEEE International Conference on Acoustics, Speech, and Signal Processing, pp 1945–1948
- Shannon CE (1948) A mathematical theory of communication. Bell System Tech J 27(3):379–423
- Sheikh HR, Bovik AC (2006) Image information and visual quality. IEEE Transactions on Image Processing 15(2):430–444
- Sheikh HR, Bovik AC, De Veciana G (2005) An information fidelity criterion for image quality assessment using natural scene statistics. IEEE Transactions on Image Processing 14(12):2117–2128
- Sheikh HR, Wang Z, Bovik AC, Cormack L (2006) Image and video quality assessment research at LIVE. [Online]. Available: <http://live.ece.utexas.edu/research/quality/>.
- Shi W, Caballero J, Huszár F, Totz J, Aitken AP, Bishop R, Rueckert D, Wang Z (2016) Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In: IEEE Conference on Computer Vision and Pattern Recognition, pp 1874–1883
- Simoncelli EP, Adelson EH (1996) Noise removal via bayesian wavelet coring. In: IEEE International Conference on Image Processing, IEEE, vol 1, pp 379–382
- Simoncelli EP, Freeman WT, Adelson EH, Heeger DJ (1992) Shiftable multiscale transforms. IEEE Transactions on Information Theory 38(2):587–607
- Simonyan K, Zisserman A (2015) Very deep convolutional networks for large-scale image recognition. In: International Conference on Learning Representations, pp 1–14
- Snell J, Ridgeway K, Liao R, Roads BD, Mozer MC, Zemel RS (2017) Learning to generate images with perceptual similarity metrics. In: IEEE Transactions on Image Processing, pp 4277–4281
- Sun J, Xu Z, Shum HY (2008) Image super-resolution using gradient profile prior. In: IEEE Conference on Computer Vision and Pattern Recognition, IEEE, pp 1–8
- Szegedy C, Zaremba W, Sutskever I, Bruna J, Erhan D, Goodfellow I, Fergus R (2013) Intriguing properties of neural networks. International Conference on Learning Representations pp 1–10
- Tao X, Gao H, Shen X, Wang J, Jia J (2018) Scale-recurrent network for deep image deblurring. In: IEEE Conference on Computer Vision and Pattern Recognition, pp 8174–8182
- Teo PC, Heeger DJ (1994) Perceptual image distortion. In: Human Vision, Visual Processing, and Digital Display V, vol 2179, pp 127–141
- Tian S, Zhang L, Morin L, Déforges O (2018) A benchmark of DIBR synthesized view quality assessment metrics on a new database for immersive media applications. IEEE Transactions on Multimedia 21(5):1235–1247
- Timofte R, De Smet V, Van Gool L (2013) Anchored neighborhood regression for fast example-based super-resolution. In: IEEE International Conference on Computer Vision, pp 1920–1927
- Timofte R, Agustsson E, Van Gool L, Yang MH, Zhang L (2017) NTIRE 2017 challenge on single image super-resolution: Methods and results. In: IEEE Conference on Computer Vision and Pattern Recognition, pp 114–125

- Tomasi C, Manduchi R (1998) Bilateral filtering for gray and color images. In: IEEE International Conference on Computer Vision, IEEE, pp 839–846
- Ulyanov D, Vedaldi A, Lempitsky V (2018) Deep image prior. In: IEEE Conference on Computer Vision and Pattern Recognition, pp 9446–9454
- Vukadinovic V, Karlsson G (2009) Trade-offs in bit-rate allocation for wireless video streaming. *IEEE Transactions on Multimedia* 11(6):1105–1113
- Wang S, Rehman A, Wang Z, Ma S, Gao W (2011) SSIM-motivated rate-distortion optimization for video coding. *IEEE Transactions on Circuits and Systems for Video Technology* 22(4):516–529
- Wang S, Ma K, Yeganeh H, Wang Z, Lin W (2015) A patch-structure representation method for quality assessment of contrast changed images. *IEEE Signal Processing Letters* 22(12):2387–2390
- Wang T, Zhang L, Jia H, Li B, Shu H (2016) Multiscale contrast similarity deviation: An effective and efficient index for perceptual image quality assessment. *Signal Processing: Image Communication* 45:1–9
- Wang X, Yu K, Wu S, Gu J, Liu Y, Dong C, Qiao Y, Loy CC (2018) ESRGAN: Enhanced super-resolution generative adversarial networks. In: European Conference on Computer Vision Workshops
- Wang Z (2011) Applications of objective image quality assessment methods. *IEEE Signal Processing Magazine* 28(6):137–142
- Wang Z, Bovik AC (2011) Reduced- and no-reference image quality assessment: The natural scene statistic model approach. *IEEE Signal Processing Magazine* 28(6):29–40
- Wang Z, Simoncelli EP (2005) Translation insensitive image similarity in complex wavelet domain. In: IEEE International Conference on Acoustics, Speech, and Signal Processing, pp 573–576
- Wang Z, Simoncelli EP (2008) Maximum differentiation (MAD) competition: A methodology for comparing computational models of perceptual quantities. *Journal of Vision* 8(12):1–13
- Wang Z, Simoncelli EP, Bovik AC (2003) Multiscale structural similarity for image quality assessment. In: IEEE Asilomar Conference on Signals, System and Computers, pp 1398–1402
- Wang Z, Bovik AC, Sheikh HR, Simoncelli EP (2004) Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing* 13(4):600–612
- Watson A (1993) DCTune: A technique for visual optimization of DCT quantization matrices for individual images. *Society for Information Display Digest of Technical Papers XXIV*:946–949
- Watson AB, Yang GY, Solomon JA, Villasenor J (1997) Visibility of wavelet quantization noise. *IEEE Transactions on Image Processing* 6(8):1164–1175
- Weiner N (1950) *Extrapolation, Interpolation and Smoothing of Stationary Time Series: With Engineering Applications*. MIT Press, Cambridge, MA, USA
- Xue W, Mou X, Zhang L, Feng X (2013) Perceptual fidelity aware mean squared error. In: IEEE International Conference on Computer Vision, pp 705–712
- Xue W, Zhang L, Mou X, Bovik AC (2014) Gradient magnitude similarity deviation: A highly efficient perceptual image quality index. *IEEE Transactions on Image Processing* 23(2):684–695
- Yang CY, Yang MH (2013) Fast direct super-resolution by simple functions. In: IEEE International Conference on Computer Vision, pp 561–568
- Yang J, Wright J, Huang TS, Ma Y (2010) Image super-resolution via sparse representation. *IEEE Transactions on Image Processing* 19(11):2861–2873
- Ye P, Kumar J, Doermann D (2014) Beyond human opinion scores: Blind image quality assessment based on synthetic scores. In: IEEE Conference on Computer Vision and Pattern Recognition, pp 4241–4248
- Zhang K, Zuo W, Chen Y, Meng D, Zhang L (2017) Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE Transactions on Image Processing* 26(7):3142–3155
- Zhang L, Zhang L, Mou X, Zhang D (2011) FSIM: A feature similarity index for image quality assessment. *IEEE Transactions on Image Processing* 20(8):2378–2386
- Zhang L, Shen Y, Li H (2014) VSI: A visual saliency-induced index for perceptual image quality assessment. *IEEE Transactions on Image Processing* 23(10):4270–4281
- Zhang R, Isola P, Efros AA, Shechtman E, Wang O (2018) The unreasonable effectiveness of deep features as a perceptual metric. In: IEEE Conference on Computer Vision and Pattern Recognition, pp 586–595
- Zhang W, Liu Y, Dong C, Qiao Y (2019a) RankSRGAN: Generative adversarial networks with ranker for image super-resolution. In: IEEE Conference on Computer Vision and Pattern Recognition, pp 3096–3105
- Zhang W, Ma K, Zhai G, Yang X (2019b) Learning to blindly assess image quality in the laboratory and wild. CoRR abs/1907.00516, [Online]. Available: <https://arxiv.org/abs/1907.00516>
- Zhang X, Feng X, Wang W, Xue W (2013) Edge strength similarity for image quality assessment. *IEEE Signal Processing Letters* 20(4):319–322

- Zhao H, Gallo O, Frosio I, Kautz J (2016) Loss functions for image restoration with neural networks. *IEEE Transactions on Computational Imaging* 3(1):47–57
- Zhou F, Yao R, Liu B, Qiu G (2019) Visual quality assessment for super-resolved images: database and method. *IEEE Transactions on Image Processing* 28(7):3528–3541