Image Super-Resolution via Deterministic-Stochastic Synthesis and Local Statistical Rectification

WEIFENG GE*, BINGCHEN GONG*, and YIZHOU YU, The University of Hong Kong

Fig. 1. Comparison of visual results from state-of-the-art algorithms for $\times 4$ single image superresolution. (a) RDN [Zhang et al. 2018]: State-of-the-art algorithm under PSNR/SSIM. (b) SRGAN [Ledig et al. 2017]: State-of-the-art algorithm in visual quality. (c) DTSN (Deterministic-sTochastic Synthesis Net): the proposed method. (d) High-resolution ground truth.

Single image superresolution has been a popular research topic in the last two decades and has recently received a new wave of interest due to deep neural networks. In this paper, we approach this problem from a different perspective. With respect to a downsampled low resolution image, we model a high resolution image as a combination of two components, a deterministic component and a stochastic component. The deterministic component can be recovered from the low-frequency signals in the downsampled image. The stochastic component, on the other hand, contains the signals that have little correlation with the low resolution image. We adopt two complementary methods for generating these two components. While generative adversarial networks are used for the stochastic component, deterministic component reconstruction is formulated as a regression problem solved using deep neural networks. Since the deterministic component exhibits clearer local orientations, we design novel loss functions tailored for such properties for training the deep regression network. These two methods are first applied to the entire input image to produce two distinct high-resolution images.

*Both authors contributed equally to this paper.

Authors' address: Weifeng Ge; Bingchen Gong; Yizhou Yu, The University of Hong Kong, Department of Computer Science, Hong Kong.

This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2018 Association for Computing Machinery.
0730-0301/2018/11-ART260 \$15.00
<https://doi.org/10.1145/3272127.3275060>

Afterwards, these two images are fused together using another deep neural network that also performs local statistical rectification, which tries to make the local statistics of the fused image match the same local statistics of the groundtruth image. Quantitative results and a user study indicate that the proposed method outperforms existing state-of-the-art algorithms with a clear margin.

CCS Concepts: • **Computing methodologies** → **Computational photography; Image processing; Neural networks;**

Additional Key Words and Phrases: Deep Learning, Image Superresolution, Deterministic Component, Stochastic Component, Local Correlation Matrix, Local Gram Matrix

ACM Reference Format:

Weifeng Ge, Bingchen Gong, and Yizhou Yu. 2018. Image Super-Resolution via Deterministic-Stochastic Synthesis and Local Statistical Rectification. *ACM Trans. Graph.* 37, 6, Article 260 (November 2018), 14 pages. <https://doi.org/10.1145/3272127.3275060>

1 INTRODUCTION

Single image super-resolution (SISR) aims to estimate high resolution images from low resolution inputs. It has been a popular research topic in the last two decades and has recently received a new wave of interest due to deep neural networks. Image super-resolution has been used for many different applications, including facial image improvement [Baker and Kanade 2000; Hu et al. 2010], compressed image/video enhancement [Gunturk et al. 2004; Segall et al. 2003], high dynamic range imaging [Bengtsson et al. 2012],

Fig. 2. Examples of the stochastic components generated by SRGAN [Ledig et al. 2017]. SRGAN may destroy structures while producing too many stochastic details in certain areas.

image mosaicing [Capel and Zisserman 1998] as well as satellite and aerial imaging [Yildirim and Gungor 2012].

It is well known that single image super-resolution is ill-posed because certain high frequency signals in the high resolution image is irrecoverably lost in the corresponding low resolution image due to downsampling. Thus SISR does not have a unique solution, and many slightly different high resolution images could become the same low resolution image after downsampling. Therefore, a high resolution image can be decomposed into two components, a *deterministic component* and a *stochastic component*. The deterministic component contains all the low-frequency signals as well as part of the high-frequency signals highly correlated to the low-frequency signals, which implies that the high-frequency signals in the deterministic component could still be recovered from the low-frequency signals in a downsampled image. The stochastic component, on the other hand, contains the rest of the signals that have little correlation with the low-frequency signals. It is the stochastic component that makes image super-resolution ill-posed.

According to the above analysis, the biggest challenges in natural image super-resolution is two-fold: first, how to reconstruct the deterministic components in the high resolution image from the low resolution image? second, how to hallucinate plausible stochastic components for the high resolution image which is compatible with the low resolution image?

To tackle these challenges, it is necessary to distinguish salient curvilinear structures, such as edges and contours, from the remaining pixels. Curvilinear structures exhibit clear and spatially coherent orientations while the orientation at other pixels, including those in texture regions, is more ambiguous or spatially incoherent. This distinction implies that the deterministic component plays a much more dominant role than the stochastic component around curvilinear structures while the stochastic component becomes more important at other pixels. We call pixels on curvilinear structures *structural pixels*, and the rest of the pixels *non-structural pixels*.

In this paper, we approach single image super-resolution from a different perspective and propose novel solutions to the aforementioned challenges. Our solution contains three cascaded stages. In the first stage, we perform deterministic component reconstruction and stochastic component hallucination independently. Deterministic component reconstruction is formulated as a regression problem solved using deep neural networks because high-frequency signals in the deterministic component are strongly correlated to signals in the low-resolution image. We design novel loss functions for training the deep regression network. Specifically, in addition to the mean absolute error (MAE) typically used for regression networks, we introduce a generalized gradient loss for better preserving local variations as well as an orientation loss that reconstructs local orientations. The orientation loss is particularly useful for structural pixels because the deterministic component at such pixels exhibits clear local orientations. For stochastic component hallucination, generative adversarial networks (GANs [Goodfellow et al. 2014]) are adopted to generate natural high frequency signals. Because GANs [Ledig et al. 2017] usually introduce artifacts that destroy image structures, they are better suited for non-structural pixels.

Although the above two methods are complementary and better suited for different subsets of pixels, both of them are applied to the entire input image to produce a pair of deterministic and stochastic components for every pixel. In the second stage, these two high resolution components are fused together to produce a single image, which should mostly preserve the deterministic component only at structural pixels and combine the two components at non-structural pixels. Since structural pixels are unknown at test time, we train another deep neural network to implicitly label structural pixels and complete the fusion. The input to this network consists of the original low-resolution input image and the aforementioned two high resolution images generated from the two complementary methods, respectively.

Although GANs can be used for generating natural stochastic components, the synthesis process is a global operation and is incapable of generating stochastic components with spatially varying statistical properties that accurately match the local deterministic component. Therefore, in the last stage, we rectify deviations introduced into the stochastic component in earlier stages by matching local statistics of the fused image with the same local statistics of the groundtruth high resolution image. Even though matching local statistics does not enforce uniqueness at the pixel level, it serves as an as-strong-as-possible constraint due to the ill-posedness of the problem. Inspired by deep learning based style transfer [Gatys et al. 2016] and texture synthesis [Sendik and Cohen-Or 2017], local Gram matrices and local correlation matrices are adopted as local statistics.

We have evaluated our method on a few popular benchmark datasets. Quantitative experimental results demonstrate that our deterministic component reconstruction algorithm achieves state-of-the-art PSNR and SSIM on Set5 [Bevilacqua et al. 2012], Set14 [Zeyde et al. 2010], B100 [Martin et al. 2001], Urban100 [Huang et al. 2015] and DIV2K validation set [Timofte et al. 2017]. On the other hand, the synthesized stochastic component may be either perceptually or statistically similar to the stochastic details in the ground truth,

but does not have strict pixelwise correspondences with the latter. Therefore, subjective evaluation based on user studies is better suited than quantitative measures based on pixel-level registration for evaluating image super-resolution results with stochastic components. To this end, a user study with 42 participants concludes that in comparison to existing state-of-the-art super-resolution algorithms, our proposed method is capable of generating high resolution images with significantly better visual quality.

In summary, this paper has the following contributions:

- A novel pipeline for single image super-resolution is introduced by decomposing SISR into three consecutive stages, deterministic component reconstruction and stochastic component hallucination, deterministic and stochastic component fusion, and local statistical rectification.
- A deep learning algorithm is designed for deterministic component reconstruction. Novel loss functions, including a generalized gradient loss and an orientation loss, are proposed to enforce properties of the deterministic component.
- A deep learning algorithm is designed for image fusion as well as local statistical rectification. Novel loss functions based on local Gramm matrices and local correlation matrices are proposed to make local statistics of the generated high resolution image match the same local statistics of the groundtruth high resolution image.

2 RELATED WORK

2.1 Exemplar or Dictionary Based Super-Resolution

Exemplar-based methods require a database of external images, and synthesize a high-resolution version of the input image by searching for exemplars in the image database and transferring relevant patches from the retrieved exemplars [Chang et al. 2004; Freeman et al. 2002]. To synthesize high-quality results, the image database needs to be fairly large, increasing the computational cost for exemplar and patch-level retrieval. Dictionary-based methods aim to mitigate this problem by learning a compact dictionary in a feature space [Timofte et al. 2013, 2014]. Alternatively, over-complete dictionaries are learned in methods based on the sparse signal reconstruction theory [Jianchao et al. 2008; Zeyde et al. 2010]. Such methods build over-complete dictionaries for both low-resolution and high-resolution images, and make certain assumptions about the sparse coefficients for the dictionaries in both domains. Because there could be multiple high resolution images corresponding to a given low resolution image, dictionary-based methods in essence average multiple possible solutions and usually generate blurry results.

Almost all exemplar-based super-resolution algorithms need a high-resolution initial image upsampling the low-resolution input image. Edge-based methods [Fattal 2007; Zhou et al. 2011] can be used to produce better initial images. Tai *et al.* [2010] further combined edge-based methods with exemplar-based synthesis. There exist major differences between our proposed method and the algorithm in [Tai et al. 2010]. First, the deterministic component in our model exists everywhere, including texture regions, and not just at edge pixels. Second, while Tai *et al.* [2010] treat all edge pixels equally, edge pixels in texture regions are not structural pixels in our method because they do not exhibit spatial coherence to form

lines or curves. In addition to such conceptual differences, at the algorithmic level, our method is based on deep neural networks and learning while Tai *et al.* [2010] rely on exemplar-based synthesis.

2.2 Deep Learning Based Super-Resolution

Given pairs of low-resolution and high-resolution images, super-resolution can be naturally cast as a regression problem that maps low-resolution images to their corresponding high-resolution images. On the other hand, deep neural networks are well known for their strong regression capability. Therefore, in recent years, many researchers have approached image super-resolution using deep learning. SRCNN [Dong et al. 2014] is a deep network with three convolutional layers designed for super-resolution regression. Other researchers have focused on designing better network architectures and more powerful loss functions. Unlike SRCNN, VDSR [Kim et al. 2016] learns the difference between the high-resolution image and the low-resolution image with a much deeper convolutional network. SRResNet [Ledig et al. 2017] shows that multiple concatenated residual blocks can further improve the regression performance. EDSR [Lim et al. 2017] uses very deep and wide residual networks and the mean absolute error for both the original images and the image gradients. LapSRN [Lai et al. 2017a] uses both residual networks and the Laplacian pyramid to perform regression at multiple intermediate resolutions. Although using the mean absolute error or the mean squared error as the training loss favors numerical measures, such as peak signal-to-noise ratio (PSNR), such loss functions do not emphasize the visual quality of the results.

As discussed earlier, SISR is ill-posed and only the high-resolution deterministic component can be uniquely reconstructed from the low-resolution image. Although the resulting images have high-resolution structures, they lack interesting high-frequency details from the stochastic component. Therefore, regression alone is unable to produce results that closely resemble natural images. To this end, generative adversarial networks (GANs) can help to a certain extent. SRGAN [Ledig et al. 2017] incorporates a perception-friendly discriminator that is trained to tell what kind of high-frequency details look natural. High-resolution images generated from SRGAN do have higher visual quality than those from regression only networks. Nevertheless, GANs do not completely solve the problem. They can only broadly tell whether certain high-frequency details look natural, but cannot tell whether such details are consistent with the given low-resolution input, or more specifically, consistent with every local region in the low-resolution input. This is the reason why SRGAN often generates high-resolution images where certain high-frequency details do not appear to be in the right place. In addition, GAN sometimes generates naturally looking high-frequency details at the cost of destroying existing image structures in the low-resolution input.

In contrast, the method proposed in this paper handles structural pixels and non-structural pixels in a more balanced way. To preserve image structures, high-resolution structural pixels are still reconstructed using regression but with additional novel loss functions. The synthesized stochastic component not only looks natural, but also is consistent with local contents in the low-resolution input.

2.3 Texture Enhancement Based Super Resolution

Previous work exists on generating high-resolution textures for SISR. Huang *et al.* [2015] exploit self-similarity of an image at different scales and adopt a patch-based transformation model to handle geometric transformations. Ahn *et al.* [2016] first generate a high-resolution texture and then use this texture image to produce another high-resolution style image with scaling and tiling. Finally they apply a style transfer algorithm to transfer the high-resolution style to the target image. Sun *et al.* [2017] complete main image structures with user interactions and then generate textures via structure propagation. They treat SISR as a constrained texture transfer problem, and need an example texture/style image when hallucinating high-frequency image details.

High-frequency details in natural images are usually complex and spatially varying. Therefore, existing SISR methods based on texture enhancement are not well suited for the intended task because they either use homogeneous textures or fail to specify the spatial correspondences between the texture style image and the estimated high-resolution image. In our proposed method, we aim to synthesize high-resolution stochastic details that satisfy spatially varying statistical properties defined by local Gram matrices and local correlation matrices.

2.4 Style Transfer and Texture Synthesis.

Gatys *et al.* [2016] use global Gram matrices computed for a subset of feature maps in a convolutional neural network to represent the style of an image. They transfer the style from a style image to a target image by minimizing a content loss and a style loss. There has been much work that adopts the same style representation while improving various other aspects, such as efficiency. Sendik *et al.* [2017] use a global correlation matrix computed for the top convolutional layer in a deep neural network to perform texture synthesis and achieve state-of-the-art results. In contrast, to be able to rectify high-frequency details with spatially varying statistics, we generalize the global Gram matrix and global correlation matrix used in such work to densely computed local Gram matrices and local correlation matrices during local statistical rectification of high-resolution stochastic components.

3 OVERVIEW

Given a low-resolution image as input, SISR aims to estimate a high-resolution counterpart that meets the following two requirements: 1) it looks like a natural image, and 2) its downsampled version is the same as the low-resolution input. For SISR based on supervised learning, one training image pair contains a high-resolution image I^{hr} and its downsampled low-resolution version (usually bicubic interpolation) I^{lr} . In our approach, SISR is decomposed into three sequential stages: the first stage performs high-resolution reconstruction of the deterministic component $I^{hr,d}$ and high-resolution hallucination of the stochastic component $I^{hr,t}$ independently; the second stage fuses the high-resolution deterministic and stochastic components generated in the previous stage; in the final stage, spatially varying local statistics are introduced to rectify the stochastic component in the fused image, and produces the final high-resolution image.

4 HIGH-RESOLUTION DETERMINISTIC COMPONENT RECONSTRUCTION

4.1 Network Architecture

As shown in Fig. 3, we adopt an existing deep network, called residual dense network (RDN [Zhang et al. 2018]) for deterministic component reconstruction. This network is composed of 4 convolutional layers, 1 upscaling module, and 16 residual dense blocks (RDBs). The detailed structure inside each module is also shown in Fig. 3. Each residual dense block has local dense connections among all of its convolutional layers. It takes the features from the preceding RDB as the input, and feed these features to all of its own convolutional layers. Each convolutional layer in a residual dense block also receives features from all the preceding convolutional layers in the same RDB. In addition, RDBs adopt the global residual learning mechanism in SRResNet [Ledig et al. 2017] to combine shallow features near the input and deep features near the output together, resulting in better regression ability.

4.2 Training Loss

In addition to the conventional color loss, we design novel loss functions for training the above deep network. Specifically, we introduce a generalized gradient loss for better preserving local variations and an orientation loss that reconstructs local orientations. The generalized gradient loss suppresses noise while preserving gradients. The orientation loss preserves curvilinear structures, including lines and curves.

4.2.1 Color Loss. To make the reconstructed image match the pixel colors of the groundtruth high-resolution image I^{hr} , we use their mean absolute difference ($\mathcal{L}_1$ norm) as the color loss, which is defined as follows.

$$\mathcal{L}_{color} = \frac{1}{Z_c} \sum_{m,n} \|I_{m,n}^{hr,d} - I_{m,n}^{hr}\|_1, \quad (1)$$

where $I_{m,n}^{hr,d}$ is the color value vector at (m,n) in $I^{hr,d}$, $I_{m,n}^{hr}$ is the color value vector at (m,n) in I^{hr} , $Z_c = 3MN$, M and N are the numbers of rows and columns of I^{hr} , and 3 is the number of color channels.

4.2.2 Generalized Gradient Loss. To emphasize edges and local structures during image reconstruction, we design a generalized gradient loss. Inspired by handcrafted features, such as local binary patterns (LBP), which compare pairs of nearby pixels, a generalized gradient vector at pixel (m,n) is defined as a feature vector holding pairwise differences between $I(m,n)$ and all other pixels in its $r \times r$ neighborhood. This feature is called generalized gradient because a conventional gradient for a discrete image is defined as a vector with two pairwise differences. The generalized gradient includes more structural information because it can use a larger neighborhood and a much larger number of pixel pairs. To improve the quality of deterministic component reconstruction, we wish the generalized gradients in $I^{hr,d}$ to be as close to the corresponding generalized gradients in I^{hr} as possible.

Fig. 3. Pipeline for high resolution deterministic component reconstruction. (a) Low resolution input image I^{lr} . (b) The deterministic component reconstruction network $\phi_d(\cdot, \theta_d)$, which is composed of 4 convolutional layers, 1 concatenation layer, 1 upsampling module, and 16 residual dense blocks. (c) Training loss $\mathcal{L}_d$, which contains a mean color loss $\mathcal{L}_{color}$, a generalized gradient loss $\mathcal{L}_{ggrad}$ and an orientation loss $\mathcal{L}_{Gabor}$. (d) A residual dense block. (e) The upscaling modules for $\times 2$, $\times 3$ and $\times 4$ superresolution.

Thus, the generalized gradient loss $\mathcal{L}_{ggrad}$ is defined as follows.

$$\mathcal{L}_{ggrad} = \frac{1}{Z_g} \sum_{m,n} \sum_{i,j \in R_{mn}} \left\| \left(I_{m,n}^{hr,d} - I_{i,j}^{hr,d} \right) - \left(I_{m,n}^{hr} - I_{i,j}^{hr} \right) \right\|_1, \quad (2)$$

where $Z_g = 3MN(r^2 - 1)$, and R_{mn} is a neighborhood of the pixel at (m, n) . Neighborhood size is set to 15×15 ($r = 15$).

4.2.3 Orientation Loss. Gabor filters are widely used in image processing because of its strong ability in texture analysis. In the spatial domain, a 2D Gabor filter is essentially a Gaussian kernel modulated by a sinusoidal plane wave. By modifying the standard deviation of the Gaussian kernels in the two principal directions or changing the frequency and phase of the sinusoidal plane wave, we can obtain a set of 2D Gabor filters. Here we follow the parameter setting used in [Ahn and Nam 2016]. Our Gabor filters have 4 scales and a dense set of 18 orientations. We set up a filter every 10 degrees to accurately characterize curvilinear structures with different orientations. In addition, we adopt three different ratios between the standard deviations in the two principal directions—1 : 1, 1 : 2, and 2 : 1. When both real and imaginary parts are used, in total, there are 432 Gabor kernels. Supplementary materials show a subset of these filters.

The orientation loss $\mathcal{L}_{Gabor}$ measures the mean difference between corresponding Gabor filter responses once they have been applied to both $I^{hr,d}$ and I^{hr} , and can be written as follows.

$$\mathcal{L}_{Gabor} = \frac{1}{Z_G} \sum_{m,n} \left\| G(I^{hr,d})_{m,n} - G(I^{hr})_{m,n} \right\|_1, \quad (3)$$

where $C_G = 432$ is the number of Gabor filters, $Z_G = 3C_GMN$, and $G(\cdot)$ represents the concatenation of Gabor filter response maps of 3 image channels. In our experiments, the spatial support of a Gabor filter is a local image patch with 51×51 pixels.

4.2.4 Final Training Loss. The final training loss for deterministic component reconstruction is a weighted sum of all three previously defined losses.

$$\mathcal{L}_d = \mathcal{L}_{color} + \alpha_1 \mathcal{L}_{ggrad} + \alpha_2 \mathcal{L}_{Gabor}, \quad (4)$$

where $\alpha_1 = 1$ and $\alpha_2 = 1$ in all our experiments. These parameters are set to make the gradients of the three losses have similar magnitude.

5 HIGH-RESOLUTION STOCHASTIC COMPONENT HALLUCINATION

Generative adversarial networks [Goodfellow et al. 2014] have been successfully incorporated into recent image superresolution algorithms [Ledig et al. 2017], which are capable of generating natural high-resolution details by exploiting a combination of color loss, perceptual loss and adversarial loss in the training process. However, we have observed that the adversarial loss cannot well preserve curvilinear structures with spatially coherent orientations, as shown in Fig 2. This is due to the fact that the architecture of GANs is not tailored for modeling spatially varying long-range correlations. Because of this, we limit the application of GANs to high resolution stochastic component hallucination only. In practice, we use the SRGAN algorithm in [Ledig et al. 2017]. Note that in the loss functions for SRGAN, the color loss is for reconstructing the deterministic component, and the perceptual loss and adversarial loss are used for generating the stochastic component. As a result, results from SRGAN have both deterministic and stochastic components. To extract the stochastic component, we obtain the deterministic component from a second network trained with the color loss only, and subtract the deterministic component from the final image generated from SRGAN. The second network shares the same architecture as the generator network in SRGAN.

Fig. 4. Pipeline for image structure detection. From left to right. (a) The ground truth image $\mathcal{I}^{hr}$ and the high resolution deterministic component $\mathcal{I}^{hr,d}$. (b) Detected curves in GT image and the high resolution deterministic component. (c) Curve filtering. (d) Image structure detection by curve filtering.

6 DETERMINISTIC AND STOCHASTIC COMPONENT FUSION

We need to fuse the two complementary components generated in the previous stage to generate a single high-resolution image. In general, every pixel in the fused image should be a weighted sum of these two components. Since the deterministic component can be used for all pixels while the stochastic component from SRGAN is suited for non-structural pixels only, the fused image should take the deterministic component only at structural pixels and combine the two components at non-structural pixels. Since groundtruth structural pixels are unknown at test time, we cannot directly generate the fused image by following this observation. Fortunately, structural pixels can be explicitly detected in groundtruth high resolution images. A mask of structural pixels for every groundtruth image can be incorporated into the training data for a deep neural network, which implicitly labels structural pixels and completes the fusion at test time.

6.1 Network Architecture

As shown in Figure 5, our deep network for deterministic and stochastic component fusion, $\phi_s(\cdot, \theta_s)$, is composed of convolution layers, residual blocks and upscaling layers. Each residual block has two convolutional layers. There are 16 residual blocks in total and the nonlinear activation function is PReLU [He et al. 2015]. The input to this deep network consists of three parts: the original low-resolution input image $\mathcal{I}^{lr}$ and two high resolution images. The first high resolution image, $\mathcal{I}^{hr,d}$, is generated from the deterministic component reconstruction network, and the other high resolution image, $\mathcal{I}^{hr,t} + \mathcal{I}^{hr,d}$, is the stochastic component generated using SRGAN [Ledig et al. 2017].

6.2 Network Training

To make the above deep network produce correct image fusion results, we need to prepare a training set. Each training sample in the

training set includes the aforementioned three input images and a corresponding binary mask separating structural pixels from non-structural ones. This binary mask is only used in the loss function for training the network. The loss function evaluates the quality of the fused image produced by the network. Pixel colors at structural pixels of the fused image should match those from the first high resolution image fed to the network (the deterministic component) while pixel colors at non-structural pixels of the fused image should be the sum of those from the two high resolution images (deterministic + stochastic components). The binary mask includes all pixels in the 5×5 neighborhood of any image structures detected in the next section. Given such a training set, the guided regression loss, $\mathcal{L}_{greg}$, for training the fusion network can be defined as follows.

$$\mathcal{L}_{greg} = \sum_{m,n} \frac{\mathcal{M}_{m,n}}{3m_0} \left\| \mathcal{I}_{m,n}^{hr,s} - \left(\mathcal{I}_{m,n}^{hr,d} + \mathcal{I}_{m,n}^{hr,t} \right) \right\|_1 + \sum_{m,n} \frac{1 - \mathcal{M}_{m,n}}{3m_1} \left\| \mathcal{I}_{m,n}^{hr,s} - \mathcal{I}_{m,n}^{hr,d} \right\|_1 \quad (5)$$

where $\mathcal{I}^{hr,s}$ represents the fused image, $\mathcal{M}$ is the binary mask ($\mathcal{M}(x, y) = 0$ means the pixel at (x, y) is a structural pixel), m_c denotes the number of color channels, m_0 is the total number of structural pixels, and m_1 is the total number of non-structural pixels.

6.3 Image Structure Detection

Image structure detection aims to identify important curvilinear structures that define the content and layout of an image. Most of these structures lie along object boundaries. However, traditional edge detection produces too many unimportant intensity/color changes, especially in the presence of high frequency textures. Deep learning based methods, such as HED [Xie and Tu 2015] and COB [Maninis et al. 2018], actually perform edge detection according to changes in local semantic meaning. They cannot accurately locate edges and contours, thus cannot be used in image structure

Fig. 5. Pipeline for image fusion and local statistical rectification. (a) Deterministic component, stochastic component and low resolution input image. (b) Image fusion and local statistical rectification network. (c) The local Gram matrix loss $\mathcal{L}_{Gram}$, and the local correlation matrix loss $\mathcal{L}_{corr}$ in the combined loss $\mathcal{L}_t$. (d) A residual block. (e) The guided regression loss $\mathcal{L}_{greg}$ in the combined loss $\mathcal{L}_t$.

detection. ELSD [Pătrăucean et al. 2012] detects and vectorizes curvilinear image structures, including line segments, circular arcs and elliptic arcs, by fitting relevant conic sections to detected edges. In this paper, we identify structural pixels through the primitives vectorized by ELSD.

Since the deterministic component is reconstructed from a low resolution image, not all curvilinear structures in the groundtruth high resolution image belong to the deterministic component. Structural pixels are defined as pixels on the groundtruth curvilinear structures that can be reconstructed from the low resolution image. This definition implies a structural pixel should lie on vectorized curves in two high resolution images, the groundtruth high resolution image and the high resolution image generated from the deterministic component reconstruction network. Thus, structural pixel labeling can be achieved as follows. First, use ELSD to detect and vectorize curves in both aforementioned high resolution images. Second, identify pixels lying on both sets of vectorized curves as structural pixels.

Let $\mathcal{S} = \{s_i\}_{i=1}^{N_S}$ be the set of vectorized curves in the high resolution image, $\mathcal{I}^{hr,d}$, from the deterministic component reconstruction network. For a curve s_i in $\mathcal{S}$, its analytic expression can be written

as follows.

$$\begin{bmatrix} x & y & 1 \end{bmatrix} \begin{bmatrix} w_i^{11} & w_i^{12} & w_i^{13} \\ w_i^{21} & w_i^{22} & w_i^{23} \\ w_i^{31} & w_i^{32} & w_i^{33} \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = 0, \quad (6)$$

where $x_1 \leq x \leq x_2$, $y_1 \leq y \leq y_2$. (x_1, y_1) is the starting point, and (x_2, y_2) is the end point. $\mathbf{w}_i$ is the 3×3 parameter matrix defining the i -th curve s_i in $\mathcal{S}$.

Let $\mathcal{T} = \{t_i\}_i^{N_T}$ be the set of vectorized curves in the groundtruth high resolution image. Since curves in $\mathcal{S}$ may not necessarily appear in $\mathcal{T}$, to obtain the subset of image structures that can be reconstructed from low resolution images, we identify the subset of curves in $\mathcal{S}$ that lie closely to curves in $\mathcal{T}$. The closeness score σ_i of a curve s_i in $\mathcal{S}$, when it is compared with curves in $\mathcal{T}$, is calculated as follows.

$$f(\mathbf{z}_k; \mathbf{w}_m^t) = \min_{\mathbf{z} \in t_m} \|\mathbf{z}_k - \mathbf{z}\|^2, \quad (7)$$

$$\sigma_i = \frac{1}{K_i} \sum_{\mathbf{z}_k \in s_i} \chi \left(\delta_a - \min_m f(\mathbf{z}_k; \mathbf{w}_m^t) \right), \quad (8)$$

where $\mathbf{z}_k = [x_k, y_k, 1]$ is the homogeneous coordinate vector of the k -th point on s_i , $\mathbf{z} \in t_m$ implies the following constraint $\mathbf{z}^T \mathbf{w}_m^t \mathbf{z} = 0$ ($x_1 < x < x_2$, $y_1 < y < y_2$), (x_1, y_1) and (x_2, y_2) are the starting

and end points of t_m , and $\mathbf{w}_m^t$ is the parameter matrix of the m -th curve t_m in $\mathcal{T}$. $\chi(x)$ represents a step function. When $x > 0$, $\chi(x) = 1$; otherwise, $\chi(x) = 0$. K_i denotes the total number of points on s_i , and $\delta_a (= 4)$ is a distance threshold. If $\sigma_i > \delta_b (= 0.8)$, the curve s_i is added into the vectorized structure set $\mathcal{E} = \{e_i\}_{i=1}^{N_{\mathcal{E}}}$. Finally, pixels on or near the curves in $\mathcal{E}$ are labeled as structural pixels, and a binary mask can be generated for them.

7 LOCAL STATISTICAL RECTIFICATION

As shown in Figure 8, although the fusion step in the previous section produces results with clear structures and natural stochastic details, it can still be further enhanced. For example, a minor stochastic component could be added onto the structural pixels in the fused image. In addition, since GANs are unsupervised, they cannot make a synthesized stochastic component have prescribed local statistical properties. To make the synthesized results more natural looking and eliminate deviations introduced into the stochastic component in earlier stages, we propose a local statistical rectification step, which matches the local statistics of the synthesized image with the same local statistics of the groundtruth high resolution image. We discover that the combination of a local Gram matrix loss and a local correlation matrix loss can achieve this goal.

7.1 Local Gram Matrix

The feature map, F , generated by a convolutional layer can be computed once we run a convolutional neural network over an image I . For the location (m, n) in the feature map, in its $r \times r$ neighborhood, we calculate a local Gram matrix $G^l(m, n) \in \mathbb{R}^{C_l \times C_l}$ in a similar manner as in [Gatys et al. 2015] except that the Gram matrix is computed for a square window centered at the location instead of over the entire feature map. Here C_l is the number of channels in the feature map generated by the l -th convolutional layer. The element $G_{ij}^l(m, n)$ in the local Gram matrix represents the correlation between the i -th channel and the j -th channel of the l -th layer in the $r \times r$ neighborhood of the location.

$$G_{ij}^l(m, n) = \frac{1}{r^2} \sum_{k \in R_{mn}} F_{ik}^l F_{jk}^l, \quad (9)$$

where R_{mn} is the $r \times r$ neighborhood centered at (m, n) , and F_{ik}^l is the value of the i -th channel of the feature map generated by the l -th convolutional layer at location $k \in R_{mn}$. Let I and $\hat{I}$ be the groundtruth image and the synthesized image, and G^l and $\hat{G}^l$ be their respective collections of local Gram matrices computed for the l -th layer. Then the local Gram matrix loss (L2 norm) for the l -th layer is calculated as follows,

$$\mathcal{L}_{Gram}^l = \frac{1}{2C_l M_l N_l} \sum_{m,n} \sum_{i,j} \|G_{ij}^l(m, n) - \hat{G}_{ij}^l(m, n)\|_2. \quad (10)$$

7.2 Local Correlation Matrix

The Gram matrix describes the degree of correlation between different feature channels. Here we also incorporate spatial correlation matrixes to describe the stochastic component over local regions. For the location (m, n) in the l -th convolutional layer, the local correlation matrix $R^{l,c} \in \mathbb{R}^{M \times N}$ for a neighborhood of (m, n) and the

c -th feature channel is defined as follows,

$$R_{ij}^{l,c}(m, n) = \frac{1}{r^2} \sum_{p,q \in R_{mn}} w_{ij}^{mn} f_{p,q}^{l,c} f_{p-i,q-j}^{l,c}, \quad (11)$$

where $f_{p,q}^{l,c}$ is the value of the c -th feature channel of the l -th convolutional layer at location (p, q) , $i \in [-\frac{r}{2}, \frac{r}{2}]$ and $j \in [-\frac{r}{2}, \frac{r}{2}]$, and $(p-i, q-j) \in R_{mn}$. $R_{ij}^{l,k}(m, n)$ represents shifting the local feature patch $f_{m,n}^{l,k}$ by i locations vertically and j locations horizontally and then performing a point-wise multiplication between the original and shifted patches. w_{ij}^{mn} is the inverse of the area of the overlapping region between these two patches, and is written as

$$w_{ij}^{mn} = [(r-|i|)(r-|j|)]^{-1}. \quad (12)$$

Then the local correlation matrix loss for the l -th layer can be written as follows.

$$\mathcal{L}_{corr}^l = \frac{1}{2C_l M_l N_l} \sum_k \sum_{m,n} \sum_{i,j} \|R_{ij}^{l,k}(m, n) - \hat{R}_{ij}^{l,k}(m, n)\|_2. \quad (13)$$

7.3 Combined Loss

We use the same network architecture for both image fusion and local statistical rectification. We find out that the loss for image fusion still needs to be incorporated when local statistical rectification is performed. Thus, the loss for local statistical rectification is a combination of the loss for image fusion, the local Gram matrix loss and the local correlation matrix loss. The combined loss $\mathcal{L}_t$ is defined as follows.

$$\mathcal{L}_s = \sum_l (\beta_0^l \mathcal{L}_{Gram}^l + \beta_1^l \mathcal{L}_{corr}^l) + \mathcal{L}_{reg}, \quad (14)$$

where (β_0^l, β_1^l) are weights for the local Gram matrix loss and local correlation matrix loss respectively.

We use the VGG-16 network [Simonyan and Zisserman 2014] pre-trained on ImageNet ILSVRC 2012 training set [Russakovsky et al. 2015] to compute the feature maps used in the local Gram matrix loss and local correlation matrix loss. Denote layers "Pool1", "Pool2", "Pool3" and "Conv4_3+ReLU" as the 1-st, 2-nd, 3-rd and 4-th layer, respectively. As shown in Figure 5, the local correlation matrix loss is only defined over the feature map generated by the layer "Pool2", and the local Gram matrix loss is defined over the feature maps generated by layers "Pool3" and "Conv4_3+ReLU", respectively. β_1^2 is set to 1e-11, both β_0^3 and β_0^4 are set to 1e-10. The weights of individual losses are set to keep their gradients at the same order of magnitude. Neighborhood size is set to 7×7 , 5×5 and 3×3 for $\mathcal{L}_{corr}^2$, $\mathcal{L}_{Gram}^3$ and $\mathcal{L}_{Gram}^4$, respectively.

8 EXPERIMENTAL RESULTS

Both the high-resolution deterministic component reconstruction and image fusion networks are implemented using Caffe [Jia et al. 2014]. GAN-based high resolution stochastic component synthesis is implemented using Tensorflow on the basis of the code from Tensor-Layer¹. All experiments run on NVIDIA TITAN X (Maxwell) GPUs with 12GB memory. Network training for deterministic component reconstruction takes 1 week using 4 GPUs in parallel while training

¹<https://github.com/tensorlayer/srgan>

for image fusion and local statistical rectification takes two days on 1 GPU. During testing, it takes around 135 seconds to generate the final result when the input is 270×480 , and the output is 1080×1920 . Generating the deterministic component only takes 60 seconds.

8.1 High-Resolution Deterministic Component Reconstruction

8.1.1 Network Settings and Parameters. In the network for deterministic component reconstruction, $\phi_s(\cdot, \theta_s)$, the kernel size of all the convolutional layers except the 1×1 convolutional layer after the concatenation layer ("Concat" in Figure 3) is 3×3 . In Figure 3(b), the number of channels in Conv1 through Conv4 is 64. There are 8 convolutional layers in a residual dense block, and a total of 16 residual dense blocks in $\phi_s(\cdot, \theta_s)$. In the pixel shuffle module, the number of channels in Conv1 is 256, 576 and 256 for "2 $\times$ ", "3 $\times$ " and "4 $\times$ " superresolution, respectively. For "4 $\times$ " superresolution, the Conv2 layer has 256 channels. The final output layer Conv5 has 3 channels. The stride for all the convolutional layers is 1. We use ReLU as the nonlinear activation function following all convolutional layers.

8.1.2 Datasets and Performance Evaluation. In the NTIRE 2017 Challenge on Single Image Super Resolution, Timofte *et al.* [Timofte et al. 2017] released a high-quality (2K resolution) dataset, called DIV2K, which consists of 800 training images, 100 validation images, and 100 test images. It has become a standard benchmark for single image super resolution. In our experiments, we train our deep network on the training set of DIV2K, and test it on five standard benchmark datasets including Set5 [Bevilacqua et al. 2012], Set14 [Zeyde et al. 2010], B100 [Martin et al. 2001], Urban100 [Huang et al. 2015] and DIV2K validation set [Timofte et al. 2017]. Note that the ground truth of the DIV2K test set is not publicly available. Therefore, we use the DIV2K validation set during testing. Images in this validation set were never used for training or parameter tuning. PSNR and SSIM [Wang et al. 2004] of the Y channel (*i.e.*, luminance) of the transformed YCbCr space are used to quantitatively evaluate the performance of different SR algorithms in addition to user studies.

8.1.3 Training Details. During network training, each mini-batch contains 5 training samples. The low-resolution input in each training sample is a 36×36 RGB patch. Following the data augmentation setting in [Lim et al. 2017], every patch is flipped horizontally or vertically and then rotated by 90° randomly. We train the network with the Adam optimizer. The initial learning rate is set to 0.0001, and the weight decay is set to 0.0005. Due to large on-board memory consumption, each experiment is run in parallel on 4 TITAN X (Maxwell) GPUs. Training converges after 350000 iterations. The learning rate is decreased by half every 70000 iterations.

8.1.4 Comparisons. We have compared our high-resolution deterministic component reconstruction algorithm with six existing state-of-the-art superresolution algorithms, including SRCNN [Dong et al. 2016], LapSRN [Lai et al. 2017b], DRRN [Tai et al. 2017a], MemNet [Tai et al. 2017b], EDSR [Lim et al. 2017], and RDN [Zhang et al. 2018]. Table 1 compares PSNR/SSIM results for three scales ($\times 2$, $\times 3$, and $\times 4$) on five different benchmarks, including Set5 [Bevilacqua et al. 2012], Set14 [Zeyde et al. 2010], B100 [Martin et al. 2001],

Urban100 [Huang et al. 2015] and DIV2K validation set [Timofte et al. 2017]. All the six existing SISR algorithms participating in this comparison are deep learning based methods. They treat the SISR problem as a regression problem. The loss used their network training is the mean color loss with the L_1 norm (mean absolute error, MAE) or L_2 norm (mean squared error, MSE). From SRCNN to RDN, the performance (PSNR/SSIM) of the regression results improves due to better network architectures and training strategies. In addition to such advances, we introduce two new loss functions, the generalized gradient loss and the orientation loss, to better preserve curvilinear structures. Experiments show that our algorithm achieves the best performance on all the benchmarks for all the scales except the "2 $\times$ " scale on Set5. When compared with state-of-the-art algorithms, the degree of improvement (0.1dB) on a saturated benchmark such as DIV2K is actually substantial. In comparison to RDN [Zhang et al. 2018], which is the previous best-performing algorithm, the proposed loss functions further improve PSNR and SSIM. In particular, on the popular DIV2K validation set [Timofte et al. 2017], our combined loss function outperforms the mean color loss alone by 0.09 db, 0.10 db and 0.12 db on the "2 $\times$ ", "3 $\times$ " and "4 $\times$ " scales, respectively. Figure 6 compares visual results from different state-of-the-art algorithms. Our proposed algorithm exhibits a stronger ability than EDSR and RDN in preserving structural lines and curves.

8.1.5 Ablation Study. To verify the effectiveness of the generalized gradient loss and the orientation loss, we remove one of them from the combined loss every time, and use the rest of the terms in the loss to train the network while keeping all other settings exactly the same as before. The results for the "2 $\times$ ", "3 $\times$ " and "4 $\times$ " scales are also reported in Table 1. When we remove the orientation loss, the performance drop ranges from 0.01 db to 0.05 db on the five benchmarks. When we remove the generalized gradient loss, the performance also drops similarly. This indicates that both the generalized gradient loss and the orientation loss can improve the regression accuracy. When they are combined, a larger performance gain can be achieved.

8.2 High-Resolution Stochastic Component Hallucination

8.2.1 Network Settings and Parameters. We use the same network settings as in the original paper [Ledig et al. 2017] in our implementation. There are 16 residual blocks in the generator network. The kernel size of all the convolutional layers except the input and output layers is 3×3 . The number of channels in all the convolutional layers except the output layer is 64. The kernel size of the input and output layers is 9×9 and the number of channels in the output layer is 3. The discriminator network is a VGG-type network that has 8 convolutional layers including the input layer. The kernel size of all the convolutional layers is 3×3 . The number of channels increases by a factor of 2 every time we move from one layer to the next. Like in the VGG network, the number of channels in a convolutional layer increases from 64 to 512. "Dense 1" has 1024 outputs and "Dense 2" has 1 output.

8.2.2 Training Details. We train the network on 115 thousand images from the ImageNet dataset and the RAISE dataset. Each mini-batch has 8 training samples, and the low-resolution input in each

Table 1. Comparison of numerical results on Set5 [Bevilacqua et al. 2012], Set14 [Zeyde et al. 2010], B100 [Martin et al. 2001], Urban100 [Huang et al. 2015] and the DIV2K validation set [Timofte et al. 2017]. Average PSNR/SSIM values for scaling factors $\times 2$, $\times 3$ and $\times 4$ are reported. The best results are shown in red, and the second best are shown in blue except for the ablation study results.

Dataset	Scale	Bicubic	SRCNN	LapSRN	DRRN	MemNet	EDSR	RDN (Baseline)	w/o GGrad	w/o Orientation	DTSN(Ours)
Set5	$\times 2$	33.66/0.9299	36.66/0.9542	37.52/0.9591	37.74/0.9591	37.78/0.9597	38.11/0.9601	38.24/0.9614	38.23/0.9611	38.23/0.9611	38.23/0.9611
	$\times 3$	30.39/0.8682	32.75/0.9090	33.82/0.9227	34.03/0.9244	34.09/0.9248	34.65/0.9282	34.71/0.9296	34.72/0.9298	34.73/0.9299	34.73/0.9299
	$\times 4$	28.42/0.8104	30.48/0.8628	31.54/0.8855	31.68/0.8888	31.74/0.8893	32.46/0.8968	32.47/0.8990	32.55/0.8988	32.57/0.8995	32.58/0.9001
Set14	$\times 2$	30.24/0.8688	32.45/0.9067	33.23/0.9130	33.23/0.9136	33.28/0.9142	33.92/0.9195	34.01/0.9212	34.10/0.9216	34.09/0.9216	34.10/0.9217
	$\times 3$	27.55/0.7742	29.30/0.8215	29.79/0.8320	29.96/0.8349	30.00/0.8350	30.52/0.8462	30.57/0.8468	30.59/0.8469	30.58/0.8469	30.59/0.8469
	$\times 4$	26.00/0.7027	27.50/0.7513	28.19/0.7720	28.21/0.7721	28.26/0.7723	28.80/0.7876	28.81/0.7871	28.87/0.7878	28.84/0.7875	28.89/0.7882
B100	$\times 2$	29.56/0.8431	31.80/0.8950	31.80/0.8950	32.05/0.8973	32.08/0.8978	32.32/0.9013	32.34/0.9017	32.36/0.9020	32.36/0.9020	32.36/0.9020
	$\times 3$	27.21/0.7385	28.82/0.7973	28.82/0.7973	28.95/0.8004	28.96/0.8001	29.25/0.8093	29.26/0.8093	29.29/0.8095	29.29/0.8095	29.29/0.8095
	$\times 4$	25.96/0.6675	27.32/0.7280	27.32/0.7280	27.38/0.7284	27.40/0.7281	27.71/0.7420	27.72/0.7419	27.73/0.7421	27.73/0.7421	27.74/0.7421
Urban100	$\times 2$	26.88/0.8403	30.41/0.9101	30.41/0.9101	31.23/0.9188	31.31/0.9195	32.89/0.9353	32.89/0.9351	32.89/0.9356	32.96/0.9356	32.96/0.9356
	$\times 3$	23.14/0.6577	25.21/0.7553	25.21/0.7553	25.44/0.7638	25.50/0.7630	26.61/0.8028	26.61/0.8028	26.66/0.8038	26.65/0.8035	26.70/0.8040
	$\times 4$	23.14/0.6577	25.21/0.7553	25.21/0.7553	25.44/0.7638	25.50/0.7630	26.61/0.8028	26.61/0.8028	26.66/0.8038	26.65/0.8035	26.70/0.8040
DIV2K	$\times 2$	31.35/0.9076	-/-	-/-	-/-	-/-	35.03/0.9475	35.17/0.9478	35.26/0.9483	35.26/0.9483	35.26/0.9483
	$\times 3$	28.49/0.8339	-/-	-/-	-/-	-/-	31.26/0.8910	31.39/0.8919	31.49/0.8931	31.48/0.8928	31.49/0.8930
	$\times 4$	26.92/0.7774	-/-	-/-	-/-	-/-	29.25/0.8440	29.34/0.8442	29.44/0.8446	29.43/0.8443	29.46/0.8451

Fig. 6. Visual comparison of deterministic component reconstruction results for the $\times 4$ scale. The input images are taken from the DIV2K [Timofte et al. 2017] validation set.

Fig. 7. Effects of individual losses on deterministic component reconstruction for the $\times 4$ scale. $\mathcal{L}_{color}$ represents the results from RDN [Zhang et al. 2018].

sample is a 48×48 RGB patch. We train the network with the Adam optimizer ($\beta = 0.9$). The initial learning rate is set to 0.0001. Experiments related to this network run on a single TITAN X (Maxwell) GPU. Training converges after 1000000 iterations. The learning rate is decreased by a factor of ten every 400000 iterations. We first

train the generator alone until convergence, and then alternate the training of the generator and discriminator.

8.3 Image Fusion and Local Statistical Rectification

8.3.1 Network Settings and Parameters. There are three inputs to the image fusion and local statistics rectification network shown

Fig. 8. Visual comparison of intermediate results from our complete image superresolution pipeline for the "4x" scale.

in Figure 5. All convolutional layers except those in the upscaling module, the residual blocks and the final output layer, use 3×3 kernels and have 64 channels. The upscaling module is the same as that in $\phi_s(\cdot, \theta_s)$. In a residual block, the Conv1 layer has 64 channels and uses 1×1 kernels; the Conv2 layer has 64 channels and uses 3×3 kernels. The output layer Conv12 has 3 channels. The stride is set to 1 for all convolutional layers. PReLU is used as the activation function following all convolutional layers.

8.3.2 Datasets and Training Details. We train the image fusion and local statistics rectification network on the DIV2K training set and test its performance on the DIV2K validation set as well as other benchmarks. We train the image fusion network $\phi_t(\cdot, \theta_t)$ using the guided regression loss $\mathcal{L}_{greg}$ at first. The initial learning rate is set to 0.0001, and the weight decay is set to 0.0005. The network is trained using a single TITAN X (Maxwell) GPU. Training stops after 90000 iterations. The learning rate is decreased by half every 30000 iterations. Then we continue to train the local statistics rectification network using the combined loss function $\mathcal{L}_t$. The learning rate is kept at 0.00001 at the beginning, and is decreased by half every 10000 iterations. The training stops after 30000 iterations. The Adam optimizer is used as in the previous stages.

8.3.3 Results. Figure 8 shows a visual comparison of intermediate results in our complete image superresolution pipeline for the "4x" scale. Figure 8(a) shows the high resolution deterministic component reconstruction result. Some non-structural pixels are blurred, but high-resolution edges are well reconstructed. Figure 8(b) shows GAN-based stochastic component synthesis result. GAN clearly

adds more realistic stochastic details, but gives rise to edge aliasing at the same time. Figure 8(c) shows detected image structures. Figure 8(d) indicates that reconstructed deterministic components and synthesized stochastic components can be effectively fused together using the guided regression loss. After comparing the images in Figure 8(d) and (e), we can find that local statistics rectification effectively enforces local statistical properties by refining previously synthesized stochastic components to make the fused results have more natural and realistic details. Compared with the ground truth in Figure 8(f), Figure 8(e) presents very competitive results.

In Figure 9, we compare our final superresolution results (DTSN) with those generated from state-of-the-art SISr algorithms (RDN [Zhang et al. 2018] and SRGAN [Ledig et al. 2017]). Each odd numbered row shows the raw results from different algorithms. Magnified local image patches from such results are shown in the row immediately below. These results show that RDN can produce high-quality edges and curves, but fail to generate fine stochastic details for non-structural pixels. SRGAN can generate natural stochastic details for all pixels without maintaining a proper balance between the stochastic and deterministic components. The results from the proposed superresolution algorithm show abundant stochastic details, which look more natural and pleasing than SRGAN, while exhibiting high-quality curvilinear structures, such as curves and lines. Figure 9 demonstrates that the proposed DTSN produces more natural and accurate high-resolution images than both EDSR and SRGAN, and these images are appealing even when compared with the ground truth.

8.4 User Study

As discussed earlier, with respect to a downsampled low resolution image, a high resolution image can be decomposed into two components, a deterministic component and a stochastic component. The deterministic component can be recovered from the low-frequency signals in the downsampled image by exploiting the correlation between the two. The stochastic component, on the other hand, contains the rest of the signals in the high resolution image that have little correlation with the low resolution image. By using the low resolution image as an anchor, the reconstructed high resolution deterministic component is fully registered with the groundtruth high resolution image. Therefore, any quality measures that rely on pixelwise correspondences, such as PSNR, could be used to evaluate the reconstructed results.

On the other hand, the high resolution stochastic component cannot be uniquely recovered from the low resolution image. We can only "hallucinate" this stochastic component through various image synthesis techniques, including GANs and statistical property matching. The synthesized stochastic component may be either statistically or perceptually similar to the stochastic details in the ground truth, but does not have strict pixelwise correspondences with the latter. Therefore, quality measures based on pixel-level registration is not well suited for evaluating image superresolution results with stochastic components. Instead, subjective evaluation based on user studies is a much more appropriate means to evaluate our final superresolution results.

To test the perceptual quality of different methods, we developed an online user study, which asks participants to mark superresolution images generated by SRGAN, RDN, EDSR and our method. We cropped 25 images from the DIV2K validation set in order to make sure that the group of output images from all methods can be displayed on the same screen and they contain recognizable foreground objects. All generated images are 4x larger than the input and divided into 25 groups. Each group has 4 images from different methods using the same input. Thus there are a total of 25×4 images.

Every participant is asked to assign integral marks to all images within 10 randomly chosen groups. The minimum score is 1, which corresponds to very poor perceptual quality. And the maximum score is 10, which corresponds to very high perceptual quality. The display order of 4 images within the same group is randomized for every group and participant. There were 42 worldwide participants and 6 of them claimed they were photographers. A brief tutorial was given by the online study before they started the test. We encouraged participants to spend at least 3 minutes on each group.

At the end, for each method, we collected 420 scores. The mean score of our method ($\mu_b = 7.59$) is the highest among all 4 methods' (μ_a), a paired T-test was performed between our method and every other method with the following hypotheses:

$$H_0 : \mu_a \geq \mu_b,$$

$$H_1 : \mu_a < \mu_b.$$

Table 2. Paired T-test Results ($\alpha = 0.05$)

Method	μ_a	T	P(two-tail)
SRGAN	5.69	-15.67	7.43e-44
RDN	4.76	-23.01	2.47e-76
EDSR	4.75	-23.58	7.60e-79

Hypothesis H_1 represents the mean score of our method is higher than the compared method. We test those hypotheses against SRGAN, RDN and EDSR. The results are shown in Table 2.

Note that all P values are less than 0.05 and all T values are negative, which means we can reject H_0 and accept H_1 with statistical significance. This gives rise to the conclusion that our method is capable of generating images that is significantly perceptually better.

8.5 Limitations

Our proposed method has the following limitations. First, its performance heavily relies on the accuracy of the structure detection algorithm. ELSD can only detect line segments, circular and elliptical arcs. When there exist more complex curves, their detected location may drift around the true location. Second, We have found that the local Gram matrix loss may cause halos/ringing around certain strong edges. When the local Gram matrix loss is only added to deeper layers which are farther from the input layer, the halo/ringing effect is likely to be reduced. Third, the hallucinated details during local statistical rectification are too strong in certain images. We have experimentally found that decreasing the weight of the local Gram matrix loss and the local correlation matrix loss can potentially make the hallucinated details less noticeable.

9 CONCLUSIONS

In this paper, we have presented a new pipeline for single image superresolution. Different from previous algorithms, we model a single pixel using two complementary components – a deterministic component and a stochastic component. The SISR problem is decomposed into independent problems generating these components. Then we fuse the two components using a deep neural network that also performs local statistical rectification, which tries to make the local statistics of the fused image match the same local statistics of the groundtruth image. Quantitative results and a user study indicate that the proposed method outperforms existing state-of-the-art algorithms with a clear margin.

ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers for their valuable comments. This work was partially supported by Hong Kong Research Grants Council under General Research Funds (HKU17209714).

REFERENCES

- Il Jun Ahn and Woo Hyun Nam. 2016. Texture Enhancement via High-Resolution Style Transfer for Single-Image Super-Resolution. *arXiv preprint arXiv:1612.00085* (2016).
- S. Baker and T. Kanade. 2000. Hallucinating faces. In *4th IEEE International Conference on Automatic Face and Gesture Recognition*. 83–88.
- T. Bengtsson, IY-H. Gu, M. Viberg, and K. Lindstrom. 2012. Regularized optimization for joint super-resolution and high dynamic range image reconstruction in a perceptually uniform domain. In *IEEE International Conference on Acoustics, Speech and Signal Processing*. 1097–1100.

- Marco Bevilacqua, Aline Roumy, Christine Guillemot, and Marie Line Alberi-Morel. 2012. Low-complexity single-image super-resolution based on nonnegative neighbor embedding. (2012).
- D. Capel and A. Zisserman. 1998. Automated mosaicing with superresolution zoom. In *IEEE Conference on Computer Vision and Pattern Recognition*. 885–891.
- Hong Chang, Dit-Yan Yeung, and Yimin Xiong. 2004. Super-resolution through neighbor embedding. In *Computer Vision and Pattern Recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE Computer Society Conference on*, Vol. 1. IEEE, I–I.
- Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. 2014. Learning a deep convolutional network for image super-resolution. In *European Conference on Computer Vision*. Springer, 184–199.
- Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. 2016. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence* 38, 2 (2016), 295–307.
- Raanan Fattal. 2007. Image Upsampling via Imposed Edge Statistics. In *ACM SIGGRAPH 2007 Papers (SIGGRAPH '07)*. ACM, New York, NY, USA, Article 95. <https://doi.org/10.1145/1275808.1276496>
- William T Freeman, Thouis R Jones, and Egon C Pasztor. 2002. Example-based super-resolution. *IEEE Computer graphics and Applications* 22, 2 (2002), 56–65.
- Leon Gatys, Alexander S Ecker, and Matthias Bethge. 2015. Texture synthesis using convolutional neural networks. In *Advances in Neural Information Processing Systems*. 262–270.
- Leon A Gatys, Alexander S Ecker, and Matthias Bethge. 2016. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2414–2423.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in neural information processing systems*. 2672–2680.
- B.K. Gunturk, Y. Altunbasak, and R.M. Mersereau. 2004. Superresolution reconstruction of compressed video using transform domain statistics. *IEEE Transactions on Image Processing* 13, 1 (2004), 33–43.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*. 1026–1034.
- Y. Hu, K.M. Lam, G. Qiu, and T. Shen. 2010. From local pixel structure to global image super-resolution: a new face hallucination framework. *IEEE Transactions on Image Processing* 20, 2 (2010), 433–445.
- Jia-Bin Huang, Abhishek Singh, and Narendra Ahuja. 2015. Single image super-resolution from transformed self-exemplars. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 5197–5206.
- Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. 2014. Caffe: Convolutional Architecture for Fast Feature Embedding. *arXiv preprint arXiv:1408.5093* (2014).
- Yang Jianchao, John Wright, Thomas Huang, and Yi Ma. 2008. Image super-resolution as sparse representation of raw image patches. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*. 1–8.
- Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. 2016. Accurate image super-resolution using very deep convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1646–1654.
- Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, and Ming-Hsuan Yang. 2017a. Deep Laplacian Pyramid Networks for Fast and Accurate Super-Resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, and Ming-Hsuan Yang. 2017b. Deep laplacian pyramid networks for fast and accurate super-resolution. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.* 624–632.
- Christian Ledig, Lucas Theis, Ferenc Huszar, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. 2017. Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. 2017. Enhanced deep residual networks for single image super-resolution. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, Vol. 1. 3.
- Kevis-Kokitsi Maninis, Jordi Pont-Tuset, Pablo Arbeláez, and Luc Van Gool. 2018. Convolutional oriented boundaries: From image segmentation to high-level tasks. *IEEE transactions on pattern analysis and machine intelligence* 40, 4 (2018), 819–833.
- David Martin, Charless Fowlkes, Doron Tal, and Jitendra Malik. 2001. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Computer Vision, 2001. ICCV 2001. Proceedings. Eighth IEEE International Conference on*, Vol. 2. IEEE, 416–423.
- Viorica Patrăucean, Pierre Gurdjos, and Rafael Grompone Von Gioi. 2012. A parameterless line segment and elliptical arc detector with enhanced ellipse fitting. In *Computer Vision–ECCV 2012*. Springer, 572–585.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. 2015. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision* 115, 3 (2015), 211–252.
- C.A. Segall, R. Molina, and A.K. Katsaggelos. 2003. High-resolution images from low-resolution compressed video. *IEEE Signal Processing Mag.* 20, 3 (2003), 37–48.
- Omry Sendik and Daniel Cohen-Or. 2017. Deep correlations for texture synthesis. *ACM Transactions on Graphics (TOG)* 36, 5 (2017), 161.
- Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- Libin Sun and James Hays. 2017. Super-resolution Using Constrained Deep Texture Synthesis. *arXiv preprint arXiv:1701.07604* (2017).
- Ying Tai, Jian Yang, and Xiaoming Liu. 2017a. Image super-resolution via deep recursive residual network. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Vol. 1.
- Ying Tai, Jian Yang, Xiaoming Liu, and Chunyan Xu. 2017b. Memnet: A persistent memory network for image restoration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4539–4547.
- Yu-Wing Tai, Shuaicheng Liu, Michael S Brown, and Stephen Lin. 2010. Super resolution using edge prior and single image detail synthesis. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, Vol. 00. 2400–2407. <https://doi.org/10.1109/CVPR.2010.5539933>
- Radu Timofte, Eirikur Agustsson, Luc Van Gool, Ming-Hsuan Yang, Lei Zhang, Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, Kyoung Mu Lee, et al. 2017. Ntire 2017 challenge on single image super-resolution: Methods and results. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2017 IEEE Conference on*. IEEE, 1110–1121.
- Radu Timofte, Vincent De Smet, and Luc Van Gool. 2013. Anchored neighborhood regression for fast example-based super-resolution. In *Proceedings of the IEEE International Conference on Computer Vision*. 1920–1927.
- Radu Timofte, Vincent De Smet, and Luc Van Gool. 2014. A+: Adjusted anchored neighborhood regression for fast super-resolution. In *Asian Conference on Computer Vision*. Springer, 111–126.
- Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* 13, 4 (2004), 600–612.
- Saining Xie and Zhuowen Tu. 2015. Holistically-nested edge detection. In *Proceedings of the IEEE international conference on computer vision*. 1395–1403.
- D. Yldrm and O. Gungor. 2012. A novel image fusion method using IKONOS satellite images. *Journal of Geodesy and Geoinformation* 1, 1 (2012), 27–34.
- Roman Zeyde, Michael Elad, and Matan Protter. 2010. On single image scale-up using sparse-representations. In *International conference on curves and surfaces*. Springer, 711–730.
- Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. 2018. Residual Dense Network for Image Super-Resolution. In *CVPR*.
- Qiang Zhou, Shifeng Chen, Jianzhuang Liu, and Xiaoou Tang. 2011. Edge-preserving Single Image Super-resolution. In *Proceedings of the 19th ACM International Conference on Multimedia (MM '11)*. ACM, New York, NY, USA, 1037–1040. <https://doi.org/10.1145/2072298.2071932>

Fig. 9. Visual comparison of final superresolution results from our method (DTSN) and other state-of-the-art algorithms for the $\times 4$ scale. The input images are taken from the DIV2K [Timofte et al. 2017] validation set. More results are given in the supplementary materials.