

Hybrid Example-based Single Image Super-Resolution

Yang Xian¹, Xiaodong Yang², and Yingli Tian^{1,3}

¹ The Graduate Center, The City University of New York
yxian@gradcenter.cuny.edu

² NVIDIA Research
xiaodongy@nvidia.com

³ The City College, The City University of New York
ytian@ccny.cuny.edu

Abstract. Image super-resolution aims to recover a visually pleasing high resolution image from one or multiple low resolution images. It plays an essential role in a variety of real-world applications. In this paper, we propose a novel hybrid example-based single image super-resolution approach which integrates learning from both external and internal exemplars. Given an input image, a proxy image with the same resolution as the target high-resolution image is first generated from a set of externally-learned regression models. We then perform a coarse-to-fine gradient-level self-refinement on the proxy image guided by the input image. Finally, the refined high-resolution gradients are fed into a uniform energy function to recover the final output. Extensive experiments demonstrate that our framework outperforms the recent state-of-the-art single image super-resolution approaches both quantitatively and qualitatively.

1 Introduction

Image super-resolution (SR) aims to estimate a fine-resolution image from one or multiple coarse-resolution images. It is essential in numerous applications such as desktop publishing, video surveillance, remote sensing, etc. Due to limitations of low-grade physical devices and sensors, sometimes we only have low-resolution (LR) images at disposal. Under these circumstances, we rely on SR techniques for the enhancement of image resolution.

Broadly speaking, image SR tasks can be divided into two categories: multi-image SR [1–4] and single-image SR. With only one input LR image available, single image SR is a numerically ill-posed problem. The high-resolution (HR) solution is not unique since it is a many-to-one mapping. Therefore, to alleviate the inherent ambiguity, single image SR generally relies on additional assumptions or priors to finalize a visually pleasing output.

Single image SR methods can be further classified as interpolation based, reconstruction based and example learning based. Interpolation-based approaches exploit pre-defined mathematical formula to predict intermediate pixel values given a LR image. Commonly used filters, e.g., bilinear and bicubic, are simple and efficient and thus are widely used in commercial software. These methods hinge on the weak smoothness

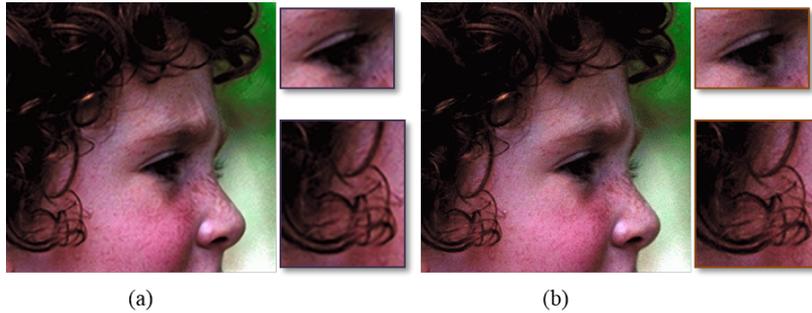


Fig. 1. SR result of image ‘face’($\times 2$). (a) HR image generated by the proposed SR method. (b) Ground truth image. Our SR approach reconstructs natural and realistic edges and textures close to the original ground truth image. This figure is better viewed on screen with HR display.

assumption and the generated HR images suffer from visual artifacts such as aliasing, jaggies, and blurring. To generate results with sharper edges, more sophisticated interpolation-based methods [5, 6] were proposed.

Reconstruction-based image upscaling methods tend to enforce some statistical priors during the estimation of the target HR image. This group of SR approaches is often referred to as “edge-directed SR” [7] due to their emphasis on restoring sharp edges. Aly and Dubois [8] incorporated a total-variation regularizer to suppress oscillatory along the edges. In [9], Shan *et al.* built a feedback-control loop that enforces the output image to be consistent with the input image. Gradient profiles [10, 11] are popularly utilized to describe the edge statistics due to its heavy-tailed distribution [12]. In [10], Fattal proposed a framework which generates the gradient field of the target image based on a statistical edge dependency relating certain edge features of two different resolutions. Sun *et al.* [11] performed a gradient field transformation to constrain the HR image gradients given the LR image based on a parametric gradient profile model. Reconstruction-based SR approaches achieve satisfying results in constructing clean edges but are less effective in hallucinating rich texture regions. A uniform parametric model for SR task is challenging since it is difficult to describe the diverse characteristics of natural images with a limited number of parameters.

Example-based SR explores the relationship between HR and their corresponding LR exemplars. The learning can be performed either via an external dataset [13–21], within the input image [22–24], or combined [25]. We refer to them as external, internal, and hybrid example-based learning. External example-based SR is brought up by Freeman *et al.* [13, 14]. Based on Freeman’s framework, many external example-based approaches are proposed to improve the SR performance and the computational speed. Coupled LR/HR dictionaries are popular representations for the raw patch exemplars or patch-related features. Yang *et al.* [15, 16] learnt a compact dictionary based on sparse signal representation which allows it to adaptively choose the most relevant reconstruction neighbors. Timofte *et al.* [19] combined the benefits of neighbor embedding and sparse coding to improve the execution speed. In [20], Zhu *et al.* proposed a deformable-patch based method by viewing a patch as a flexible deformation flow. This leads to a more “expressive” representation without increasing the size of the dictionary. Yang *et*

al. [21] divided the input feature space of LR source images into subspaces and trained a regression model for each individual subspace.

Internal example-based SR is based on the fact that small patches in a natural image tend to appear repeatedly within the image itself and across different scales. Glasner *et al.* [22] performed the nearest-neighbor search based on a patch pool formed with patches collected through a pyramid structure with the input image at different resolutions. In [23], Freedman and Fattal proposed a real time multi-step algorithm which allows a local search instead of a global search. Huang *et al.* [24] extended internal example-based SR by allowing geometric variations to fully utilize the internal similarity. For hybrid example-based learning, Yang *et al.* [25] combined learning from self-examples and an external dataset into a regression model based on in-place examples.

Effectiveness of internal example-based SR has been demonstrated in [26]. The difficulty in estimating missing high frequency details increases as the scaling factors get larger due to the increment of LR/HR ambiguity. External example-based learning breaks the limitation by introducing new information from a natural image dataset. However, the performance of external example-based SR depends on the similarity between the training dataset and the testing images. Due to diversification in natural images, the lack of relevance between certain testing images and a universal training dataset still exists. Keeping increasing the size of the training dataset provides a limited solution but still leaves the key problem untouched.

In this paper, we propose a novel hybrid example-based single image SR method which incorporates learning image-level statistics from an external dataset and gradient-level self-awareness with internal statistics. The proposed SR scheme consists of three steps: a proxy HR image is constructed through a set of pre-built regression models learnt from external exemplars; the gradients of the proxy image are then fed into a pyramid self-awareness framework guided by the input LR gradients; finally, the refined HR gradients and the input image are integrated into a uniform cost function to recover the final HR image. Fig. 1 illustrates the comparison of the generated HR image ‘face’ and the ground truth (GT) image. Our SR result is very close to the original GT image. Edge details including eye and face contours are natural and realistic. Hair textures are well reconstructed with minimal visual artifacts. The contributions of the proposed SR framework are fourfold:

- A novel single image SR scheme is proposed to benefit from both external and internal statistics. Our framework hinges on image-level hallucination from externally learnt regression models as well as gradient level pyramid self-awareness for edges and textures refinement. A uniform energy function is utilized to restore the final HR image in a manner consistent with the input image.
- In the training of a set of regression models from an external dataset with exemplar patches, we model the input LR feature space with Gaussian Mixture Model (GMM) to ensure the effective and targeted learning.
- To obtain quality edges and textures, a novel gradient-level self-refinement pyramid is proposed to recover the high-frequency details lost during the reconstruction process of proxy HR image.
- The proposed framework is effective and outperforms the recent state-of-the-art single image SR algorithms quantitatively and qualitatively.

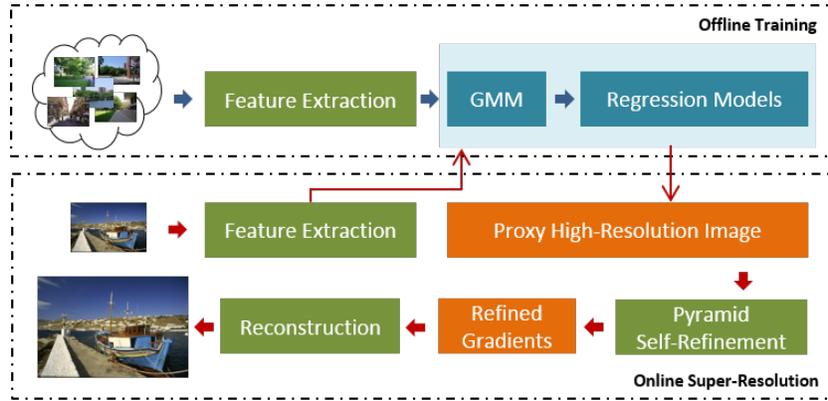


Fig. 2. Flowchart of the proposed SR method. Given a LR image, a proxy HR image is constructed through the regression models trained via an external dataset. The input feature space is modeled with GMM to ensure a targeted learning. With the proxy image, its gradients are refined using gradients of the input image. The refined HR gradients are then integrated into the reconstruction framework to recover the final output image.

2 Hybrid Example-based Super-Resolution

External example learning-based SR usually relies on learning priors or models from a natural image dataset which leads to a stable SR performance. Different from internal example-based approaches, learning externally is normally performed off-line and is less time consuming when upscaling a testing image. However, natural images vary dramatically especially for edges and textured regions. Given a natural image, certain patches occur rarely in the universal training dataset and this results in a less effective SR performance for those patterns. On the other hand, internal patch redundancy has been validated to be effective both in “expressiveness” (how similar between a small patch and its most similar patches found internally or externally) and “predictive power” (how well the found similar patches can be used in image restoration tasks given a prediction model) [26]. In order to combine the benefits of external and internal example-based learning, we propose a hybrid learning based SR framework. Fig. 2 illustrates the schematic pipeline of our approach. The system consists of three steps to upsample an image: proxy image recovery from external statistics, gradient-level self-awareness from internal statistics, and final image reconstruction.

Given a LR image, a proxy HR image is first generated with a group of pre-built regression models. The regression models are trained on an external natural image dataset. To ensure a targeted learning, the input feature space is modeled with GMM where an individual regression model is trained for each Gaussian component. The generated proxy HR image is robust with stable SR performance since the regression models are trained through a large number of natural images in a divide-and-conquer manner. However, certain LR patches in the input image may appear rarely within the training dataset and thus lead to an inaccurate HR prediction, i.e., over-smoothed with missing high-frequency details. Therefore, after obtaining the proxy image, a gradient-

level coarse-to-fine self-refinement is performed guided by gradients of the input LR image. Motivated by reconstruction-based SR approaches, we adopt a gradient-level refinement to better preserve the intensity changes. This process aims to replace the high-variance gradient patches in the proxy image with more accurate representations to recover more visually plausible outputs. Finally, the targeted HR image is restored through minimizing a uniform cost function with the refined gradients. The detailed three steps are presented in the following subsections.

2.1 Proxy Image Recovery

Given an input image L , we first recover a proxy HR image from a set of externally-trained regression models. A large set of LR/HR exemplar patch pairs with magnification factor s are collected from a dataset consisted of more than 6,000 images. All images within the dataset are considered HR images and the corresponding LR images are generated with a blur and downsampling process. To better preserve the structure information, for a LR/HR patch pair $\{P_l, P_h\}$, we normalize both patches by extracting the mean value of P_l . After normalization and vectorization, the input LR and HR features are represented as $\mathbf{X} \in \mathbb{R}^{l \times M}$ and $\mathbf{Y} \in \mathbb{R}^{r \times M}$ respectively where l and r denote the corresponding feature dimensions and M indicates the number of samples.

To ensure a targeted learning, we first model the input LR feature space where later multiple regression models are trained. The most straightforward model to describe the feature space is the normal distribution. However, a single normal distribution is insufficient to capture the complex nature of the features. We therefore employ GMM to represent the feature distribution. GMM is a generative model which has the capacity to model any given probability distribution function when the number of Gaussian components is large enough. Given a GMM with K components, the probability of a feature \mathbf{x}_i is

$$p(\mathbf{x}_i|\theta) = \sum_{k=1}^K w_k \mathcal{N}(\mathbf{x}_i; \boldsymbol{\mu}_k, \boldsymbol{\sigma}_k), \quad (1)$$

where w_k is the prior mode probability which satisfies the constraint $\sum_{k=1}^K w_k = 1$, and $\mathcal{N}(\mathbf{x}_i; \boldsymbol{\mu}_k, \boldsymbol{\sigma}_k)$ indicates the k th normal distribution with mean $\boldsymbol{\mu}_k$ and variance $\boldsymbol{\sigma}_k$:

$$\mathcal{N}(\mathbf{x}_i; \boldsymbol{\mu}_k, \boldsymbol{\sigma}_k) = \frac{\exp\left(-\frac{1}{2}(\mathbf{x}_i - \boldsymbol{\mu}_k)^T (\boldsymbol{\sigma}_k)^{-1} (\mathbf{x}_i - \boldsymbol{\mu}_k)\right)}{(2\pi)^{l/2} |\boldsymbol{\sigma}_k|^{1/2}}, \quad (2)$$

where $\mathbf{x}_i \in \mathbb{R}^l$, $\boldsymbol{\mu}_k \in \mathbb{R}^l$, and $\boldsymbol{\sigma}_k \in \mathbb{R}^{l \times l}$. By using the Expectation-Maximization (EM) algorithm to optimize the Maximum Likelihood (ML) from a large number of features, we can estimate the GMM parameters $\theta = \{w_k, \boldsymbol{\mu}_k, \boldsymbol{\sigma}_k, k = 1, \dots, K\}$. We employ 200,000 randomly sampled features to learn the parameters θ in our experiment. Though Eq. (1) supports the full covariance matrix, a diagonal matrix in practice is sufficient to model most distributions. Moreover, the GMM with diagonal matrices is more computationally efficient and stable compared to the one with full matrices.

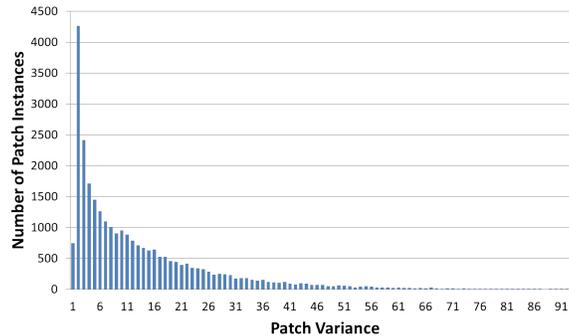


Fig. 3. Variance distributions of 20,000 patches with size 7×7 extracted from BSDS200 [27]. Given a patch, the larger the variance is, the less frequent it tends to appear within the dataset.

GMM is based on a well-defined statistical model and is computationally tractable. We then assign each LR feature $\mathbf{x}_i \in \mathbf{X}$ to corresponding Gaussian component with the highest probability. Suppose there are M_k patches associated with the k th Gaussian component and $\mathbf{X}_k \in \mathbb{R}^{l \times M_k}$, $\mathbf{Y}_k \in \mathbb{R}^{r \times M_k}$ represent the corresponding LR/HR features, a linear regression model is then trained with the regression coefficient A_k learnt through:

$$A_k^* = \operatorname{argmin}_{A_k} \{|\mathbf{Y}_k - A_k \hat{\mathbf{X}}_k|^2\}, \quad (3)$$

where $\hat{\mathbf{X}}_k = [\mathbf{X}_k^T \mathbf{1}]^T$. During testing phase, given a LR image, we first extract all features by performing normalization and vectorization for every LR patch. Then each feature is assigned to a Gaussian component according to the posterior where the corresponding regression model is applied to obtain the HR patch. We use simple averaging to blend overlapping pixels to generate the proxy HR image.

2.2 Gradient-Level Self-Awareness

Internal patch redundancy has been demonstrated powerful for image restoration tasks [26] and serves as the theoretical foundation for internal example-based SR methods [22–24]. With good performance for SR under relatively small magnification factors, the limitations of internal example-based SR approaches lie on the heavy computational costs to execute on-line exhaustive pair-wise patch comparisons and degraded performance with the increment of the scaling factor. In this step, advantage of self-similarity is absorbed to refine the proxy image generated previously without going through exhaustive patch matching.

The self refinement process aims at recovering the missing high-frequency details for patches which are not frequently seen in the external training dataset. We first verify that patches with higher variances tend to appear less frequently within a natural image dataset. The experiment is performed by randomly extracting 20,000 patches of size 7×7 within the Berkeley Segmentation Database (BSDS200) [27]. As observed from Fig. 3, the number of patch instances decreases quickly as the variance increases.

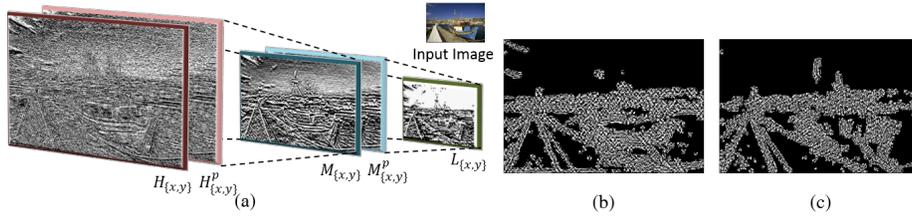


Fig. 4. (a) Illustration for gradient-level coarse-to-fine self-awareness procedure. Gradients of the proxy image $H_{\{x,y\}}^p$ are downsampled to $M_{\{x,y\}}^p$ which are refined with the gradients $L_{\{x,y\}}$ of the input image. Gradients with darker frames represent the corresponding refined results of the ones with lighter frames. Afterwards, $H_{\{x,y\}}^p$ is refined with $M_{\{x,y\}}$. Please refer to text for details. (b) Difference map between $H_{\{x\}}$ and $H_{\{x\}}^p$. (c) Difference map between $H_{\{y\}}$ and $H_{\{y\}}^p$

To refine H^p with L , patches from H^p of size $a \times a$ with variance larger than a pre-set threshold θ are firstly extracted. We utilize self-similarity to recover the missing high-frequency details of those high variance patches which are not frequently seen in the external dataset. For each high variance patch, its k most similar patches with the same size are searched and extracted within L where the similarity of two patches is measured in their mean square error (MSE). Afterwards, the original patch is replaced with the weighted sum of the found k patches in a softmax way.

Motivated by reconstruction-based SR, in our proposed scheme, a gradient-level self-refinement is adopted to better preserve edges and textures information. Moreover, it is validated in [22] that average patch recurrence across scales decays as the resolution difference increases. Therefore, if the magnification factor s is larger than s_0 ($s_0 = 3$ in our experimental setting), the proposed self-refinement is executed in a coarse-to-fine scheme.

Fig. 4 illustrates the self-awareness process. After obtaining the proxy image H^p , its gradients in horizontal and vertical (denoted as x and y) directions are computed and refined with the corresponding gradients of the input image L . In later context, for ease of interpretation, we denote the gradients of an image I in x and y directions as $I_{\{x,y\}}$. The refinement is performed separately for each gradient direction. Take direction x as an example, if s is larger than s_0 , we first downsample H_x^p by factor \sqrt{s} to obtain M_x^p . After that, high variance patches in M_x^p are refined with L_x to obtain finer-version gradient M_x . Then the final HR gradient H_x is computed by utilizing M_x to refine H_x^p . In the above process, if scaling factor \sqrt{s} is still larger than s_0 , we further decompose \sqrt{s} in a similar manner before proceed.

Gradient patches are mostly flat with small variances. Therefore, only a small portion of the patches are refined with missing high frequency details. To ensure a more effective refinement, all the patches are normalized to have zero means and unit standard variances before searching. The combined patch is then readjusted according to the original mean and variance of the input patch. After the self-awareness, proxy HR patches which are over-smoothed during the upsampling process are refined with restored high frequency details.

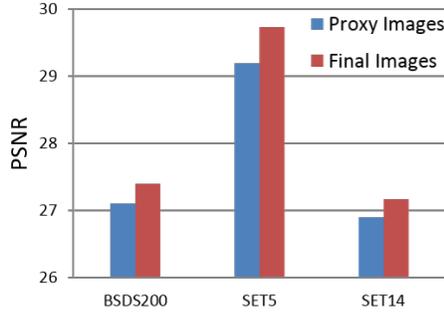


Fig. 5. Average PSNR (dB) comparisons of the proxy images (marked in blue) and the corresponding final output (marked in red) recovered from self-refinement and reconstruction in datasets BSDS200 [27], SET5 [28], and SET14 [29] (scaling factor $\times 4$). There is an obvious boost in PSNR after performing self-awareness and reconstruction for all three datasets.

2.3 Final Image Reconstruction

The final step is to reconstruct the targeted output image from the self-refined HR gradients through the following cost function:

$$H^* = \operatorname{argmin}_H \{ |\nabla H - \nabla H_r|^2 + \alpha |(H * G) \downarrow_s - L|^2 \}, \quad (4)$$

where ∇H_r represents the refined $H_{\{x,y\}}$ after the pyramid gradient-level self-awareness step. G stands for a Gaussian kernel with standard variance σ varies for different scaling factors s : $\sigma = \{0.8, 1.2, 1.6\}$ for $s = \{2, 3, 4\}$. α is the weighting factor.

Constraints in both the gradient-level and the image-level are integrated into a uniform cost function. The first term states the constraint imposed by the self-refined HR gradients. The second constraint ensures the fidelity between the output and the input images. The energy function can be optimized through the gradient descent algorithm.

We demonstrate the effectiveness of our proposed gradient-level self-refinement and the feasibility of reconstructing images based on gradients experimentally on datasets BSDS200 [27] (200 images), SET5 [28] (5 images), and SET14 [29] (14 images). All the images are downsampled by a factor of 4. Fig. 5 presents the average Peak Signal-to-Noise Ratio (PSNR) comparison of the proxy HR images and the final outputs after self-awareness and reconstruction within each dataset. After refining the ambiguous patches, all the images experience an obvious boost in the SR performance measured in PSNR.

3 Experimental Results

In this section, the proposed hybrid example-based SR method is evaluated with multiple natural images on SET5 [28], SET14 [29], and BSDS200 [27]. We also compare our results with recent state-of-the-art single image SR algorithms both quantitatively and qualitatively.

Table 1. Comparison of the proposed approach with recent state-of-the-art methods in SET5 [28], SET14 [29], and BSDS200 [27] in terms of average PSNR (dB). Our results outperform other methods in all three datasets.

Dataset	Bicubic	Shan [9]	ScSR [15]	Zeyde [29]	ANR[19]	Yang [21]	Ours
SET5[28]	28.4202	28.2488	28.1507	29.6990	29.6959	29.2540	29.7199
SET14[29]	26.1911	26.0938	26.1244	27.1149	27.0918	26.9004	27.1559
BSDS200[27]	26.6339	26.5103	26.5232	27.3286	27.3255	27.1749	27.3568

Parameter Selection: Same as many existing SR methods, for color images, we only perform the proposed SR algorithm on the luminance channel in YUV color space while the other two color channels are upsampled with bicubic interpolation.

The training dataset used for regression model learning is the same as in [21] with 6,152 natural images. We extract all patches with size 7×7 from LR images. Corners of each patch are removed and thus the LR feature dimension is 45. Only the central $3s \times 3s$ pixels in the corresponding HR patch are used to formulate the HR feature where s indicates the magnification factor. We randomly select 200,000 LR/HR features to train the GMM model with 512 components. To better model the feature space, we filter out the smooth patches before selection. With the trained GMM, each feature is assigned to a Gaussian component with the highest probability. We learn a linear regression model for each Gaussian component using maximum 1,000 LR/HR features within this component.

In the pyramid gradient-level self-awareness, the maximum magnification factor s_0 between each level is 3. If the scaling factor s is larger than 3, we adopt a coarse-to-fine scheme with a factor of \sqrt{s} per-step. Patch size a in the self-refinement is 7 and the pre-set threshold θ used to differentiate the smooth patches and the high-variance ones is 5. The number k of similar patches captured during the searching is 5. We set α in Eq. (4) to be $4/7$.

Quantitative Analysis: Our proposed approach is evaluated on a variety of natural images and we compare the generated SR results with recent state-of-the-art methods [15, 9, 29, 19, 21] quantitatively measured in PSNR. We use the source code [29, 19, 21] and executable file [9] provided by the authors or a third-party implementation [15] to generate the corresponding HR images based on the same LR input images. To be more specific, given a GT image, the LR image is obtained by performing the bicubic down-sampling. These LR images are saved in PNG format and serve as the uniform input for all SR approaches. We then follow the code or executable file provided by the authors to perform the image upsampling. SR methods [15, 29, 19] generate results with borders shaved. To perform a fair comparison, we crop the borders for all the HR results generated by different SR approaches utilizing the same scheme before the PSNR calculation over the luma channel.

The evaluation is performed on three datasets, i.e., SET5 [28], SET14 [29], and BSDS200 [27] at magnification factor 4. As illustrated in Table 1, for different input images, our proposed approach is robust and outperforms the other methods measured

in average PSNR over all three datasets.

Qualitative Analysis: Fig. 6 presents a set of SR results on image ‘Lenna’ with an upscaling factor of 4. Reconstructing the hat contours with minimal visual artifacts is difficult for most of the SR approaches listed. Our method successfully generates clear contours consistent with the GT image.

Fig. 7 provides another set of results with a scaling factor of 4 on image ‘snow’. It is challenging to reconstruct the rattan textures as illustrated by the zoom-in regions. Results generated by bicubic interpolation and [16] are over-smoothed. Deformed patterns exist within [9] (irregular squared patterns) and [21] (discontinuities). [19] fails to recover several edges and [30] over-sharps the edges. Our recovered HR image reveals natural patterns.

4 Conclusion

In this paper, we have proposed a novel hybrid example-based single image super-resolution approach which integrates both external and internal statistics. Given an input LR image, a proxy HR image is firstly generated with pre-built regression models learnt from a large external dataset. Then its gradients in horizontal and vertical directions are refined guided by the corresponding gradients of the input image. Finally, the refined HR gradients and the input LR image are fed into our proposed energy function to recover the final output.

As demonstrated by the extensive experimental results, the proposed approach is robust with satisfying super-resolution performance measured quantitatively in PSNR. The generated HR images tend to have sharper edges and more natural textures compared to recent state-of-the-art super-resolution methods.

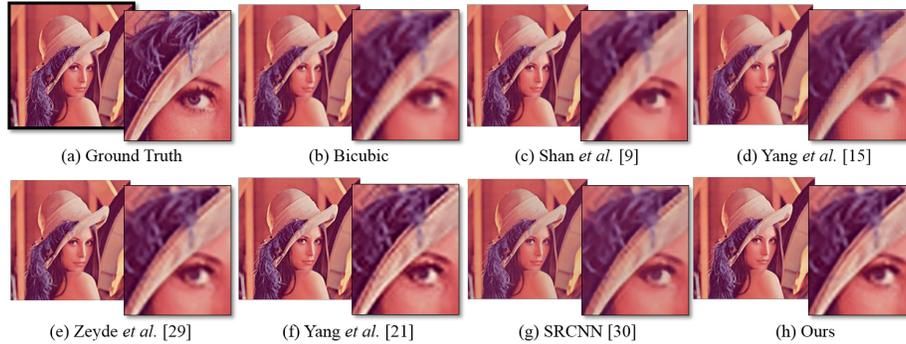


Fig. 6. SR of image ‘Lenna’ ($\times 4$). Zoom-ins clearly indicates that our proposed approach reconstructs the hat contours with minimal artifacts while other methods suffer from blurring, jaggies, and aliasing artifacts. This figure is better viewed on screen with HR display.

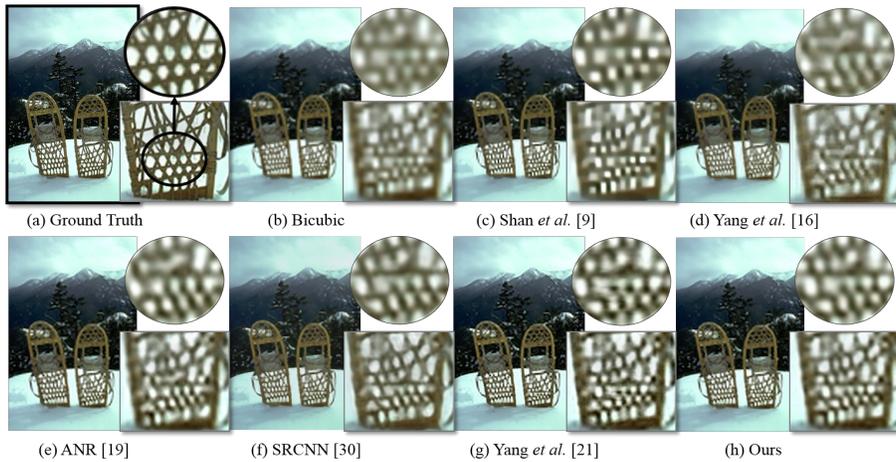


Fig. 7. SR of image ‘snow’ ($\times 4$). It is a challenging task to reconstruct the rattan textures. Results generated by bicubic interpolation and [16] over-smooth the textures. Deformed patterns exist in [9] (squared textures) and [21] (discontinuities). [19] fails to recover several edges as shown in the circled zoom-in and [30] oversharpens the edges. Our result best reconstructs the details. This figure is better viewed on screen with HR display.

Acknowledgments. This work was supported in part by ONR grant N000141310450 and NSF grants EFRI-1137172, IIP-1343402.

References

1. Boulanger, J., Kervrann, C., Bouthemy, P.: Space-time adaptation for patch-based image sequence restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **29** (2007) 1096–1102
2. Farsiu, S., Robinson, M.D., Elad, M., Milanfar, P.: Fast and robust multiframe super resolution. *IEEE Transactions on Image Processing* **13** (2004) 1327–1344
3. Protter, M., Elad, M., Tekeda, H., Milanfar, P.: Generalizing the non-local-means to super-resolution reconstruction. *IEEE Transactions on Image Processing* **18** (2009) 36–51
4. Shi, B., Ben-Ezra, H.Z.M., Yeung, S., Fernandez-Cull, C., Shepard, R.H., Barsi, C., Raskar, R.: Sub-pixel layout for super-resolution with images in the octic group. In: *ECCV*. (2014)
5. Li, X., Orchard, M.T.: New edge-directed interpolation. *IEEE Transactions on Image Processing* **10** (2001) 1521–1527
6. Su, D., Willis, P.: Image interpolation by pixel-level data-dependent triangulation. *Computer Graphics Forum* **23** (2004) 189–201
7. Tai, Y., Liu, S., Brown, M., Lin, S.: Super resolution using edge prior and single image detail synthesis. In: *CVPR*. (2010)
8. Aly, H.A., Dubois, E.: Image up-sampling using total-variation regularization with a new observation model. *IEEE Transactions on Image Processing* **14** (2005) 1647–1659
9. Shan, Q., Li, Z., Jia, J., Tang, C.K.: Fast image/video upsampling. In: *ACM SIGGRAPH Asia*. (2008)
10. Fattal, R.: Image upsampling via imposed edge statistics. In: *ACM SIGGRAPH*. (2007)

11. Sun, J., Sun, J., Xu, Z., Shum, H.Y.: Image super-resolution using gradient profile prior. In: CVPR. (2008)
12. Huang, J., Mumford, D.: Statistics of natural images and models. In: CVPR. (1999)
13. Freeman, W.T., Pasztor, E.C., Carmichael, O.T.: Learning low-level vision. *International Journal of Computer Vision* **40** (2000) 25–47
14. Freeman, W.T., Jones, T.R., Pasztor, E.C.: Example-based super-resolution. *Computer Graphics and Applications* **22** (2002) 56–65
15. Yang, J., Wright, J., Huang, T., Ma, Y.: Image super-resolution as sparse representation of raw image patches. In: CVPR. (2008)
16. Yang, J., Wright, J., Huang, T.S., Ma, Y.: Image super-resolution via sparse representation. *IEEE Transactions on Image Processing* **19** (2010) 2861–2873
17. HaCohen, Y., Fattal, R., Lischinski, D.: Image upsampling via texture hallucination. In: ICCP. (2010)
18. Sun, J., Zhu, J., Tappen, M.F.: Context-constrained hallucination for image super-resolution. In: CVPR. (2010)
19. Timofte, R., Smet, V.D., Gool, L.V.: Anchored neighborhood regression for fast example-based super-resolution. In: ICCV. (2013)
20. Zhu, Y., Zhang, Y., Yuille, A.L.: Single image super-resolution using deformable patches. In: CVPR. (2014)
21. Yang, C.Y., Yang, M.H.: Fast direct super-resolution by simple functions. In: ICCV. (2013)
22. Glasner, D., Bagon, S., Irani, M.: Super-resolution from a single image. In: ICCV. (2009)
23. Freedman, G., Fattal, R.: Image and video upscaling from local self-examples. *ACM Transactions on Graphics* **28** (2010) 1–10
24. Huang, J.B., Singh, A., Ahuja, N.: Single image super-resolution from transformed self-exemplars. In: CVPR. (2015)
25. Yang, J., Lin, Z., Cohen, S.: Fast image super-resolution based on in-place example regression. In: CVPR. (2013)
26. Zontak, M., Irani, M.: Internal statistics of a single natural image. In: CVPR. (2011)
27. Martin, D., Fowlkes, C., Tal, D., Malik, J.: A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In: ICCV. (2001)
28. Bevilacqua, M., Roumy, A., Guillemot, C., Morel, M.A.: Low-complexity single-image super-resolution based on nonnegative neighbor embedding. In: BMVC. (2012)
29. Zeyde, R., Elad, M., Protter, M.: On single image scale-up using sparse-representations. *Curves and Surfaces* **6920** (2010) 711–730
30. Dong, C., Loy, C.C., He, K., Tang, X.: Learning a deep convolutional network for image super-resolution. In: ECCV. (2014)