Super Resolution using Edge Prior and Single Image Detail Synthesis

Yu-Wing Tai[1] Shuaicheng Liu[2] Michael S. Brown[2] Stephen Lin[3]

[1]Korean Advanced Institute of Science and Technology.
[2]National University of Singapore. [3]Microsoft Research Asia.

Abstract

Edge-directed image super resolution (SR) focuses on ways to remove edge artifacts in upsampled images. Under large magnification, however, textured regions become blurred and appear homogenous, resulting in a super-resolution image that looks unnatural. Alternatively, learning-based SR approaches use a large database of exemplar images for "hallucinating" detail. The quality of the upsampled image, especially about edges, is dependent on the suitability of the training images. This paper aims to combine the benefits of edge-directed SR with those of learning-based SR. In particular, we propose an approach to extend edge-directed super-resolution to include detail from an image/texture example provided by the user (e.g., from the Internet). A significant benefit of our approach is that only a single exemplar image is required to supply the missing detail – strong edges are obtained in the SR image even if they are not present in the example image due to the combination of the edge-directed approach. In addition, we can achieve quality results at very large magnification, which is often problematic for both edge-directed and learning-based approaches.

1. Introduction and Related Work

The goal of image super resolution (SR) is to estimate a high resolution (HR) image from a low resolution (LR) input. The problem is inherently ill-posed given that many HR images can produce the same LR image when downsampled. This LR to HR image ambiguity increases as magnification becomes larger [2].

Approaches addressing the SR problem can be categorized as *interpolation based*, *reconstruction based*, and statistical or *learning based* (for a good survey see [28]). Interpolation based approaches (e.g., [1, 14, 4, 27]) have their foundations in sampling theory and try to interpolate the HR image from the LR input. These approaches usually blur high frequency details and often have noticeable aliasing artifacts along edges.

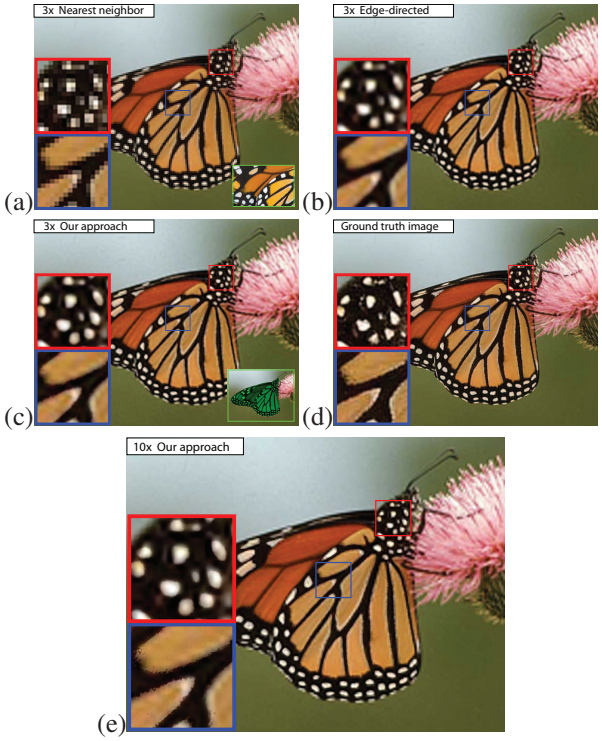

Figure 1. Example-based detail synthesis. (a) 3× magnification by nearest neighbor upsampling of an input low resolution (LR) image with a user supplied example image; (b) result using edge-directed SR [23]; (c) result from our approach that synthesizes details from the input example. The region where detail is transferred is shown in the lower right inset; (d) ground truth image; (e) 10× magnification using our approach. The example texture was found using Google image search with the keyword "monarch wing".

Reconstruction based approaches (e.g., [13, 20, 26, 17, 25, 6, 3, 23, 22]) estimate an SR image by enforcing some prior knowledge on the upsampled image (e.g., smooth edges). These approaches also require the appearance of the upsampled image to be consistent with the original LR image(s) via back-projection. The enforced priors are typically designed to reduce edge artifacts. We refer to this class of methods as edge-directed SR. The performance of reconstruction based approaches depends on the prior used

and its compatibility with the given image.

Learning based techniques estimate high frequency details from a large training set of HR images that encode the relationship between HR and LR images (e.g., [10, 19, 9, 2, 24, 5, 29, 18]). These approaches effectively "hallucinate" missing details based on similarities between the LR image and the examples in the training set. These approaches have shown great promise and have been applied to SR in various ways, including generic detail synthesis for upsampling [10, 19, 9, 2], edge-focused detail synthesis [24], imposing consistency on synthesized detail [24, 5], and targeting multiple low-resolution images [29]. As discussed in [16] the effectiveness of learning based SR methods depends highly on the supporting image database, especially as magnification increases. Recent work in [11] used a combined multi-image and learning-based strategy, where the training set is obtained from the low-resolution input itself. This particular approach is reliant on redundancies in the input image which is shown to be more beneficial in sharpening edges than in synthesizing details.

The major drawback of edge-directed SR approaches is their focus on preserving edges while leaving relatively "smooth" regions untouched. As discussed in [2, 17], if a SR algorithm targets only edge preservation, there exists a fundamental limit (about $5.5\times$ magnification) beyond which high frequency details can no longer be reconstructed. Loss of these details leads to unnatural images with large homogeneous regions. This effect is demonstrated in Figure 2 that plots the gradient statistics of SR images with different magnification factors. Shown are bicubic upsamling (b) and edge-directed SR [23] (c). The respective gradient statistics plots shown in Figure 2(d-e) increasingly deviate from the heavy-tailed distribution of natural image statistics [8] as the magnification factor increases.

To produce photo-realistic results for large magnification factors, not only must edge artifacts be suppressed, but image details lost due to limited resolution need to also be recovered. Learning based techniques can achieve the latter goal; however, as mentioned in many previous works, the performance of learning based SR depends heavily on the similarity between training data and the test images. In particular, the quality of edges in the SR image can be significantly degraded when corresponding edges in the training data do not match or align well. Accurate reconstruction of edges is critical to SR, as edges are arguably the most perceptually salient features in an image.

In this paper, we propose an approach that reconstructs edges while also recovering image details. This is accomplished by adding learning-based detail synthesis to edge-directed SR in a mutually consistent framework. Our method first reconstructs significant edges in the input image using an edge-directed super-resolution technique,

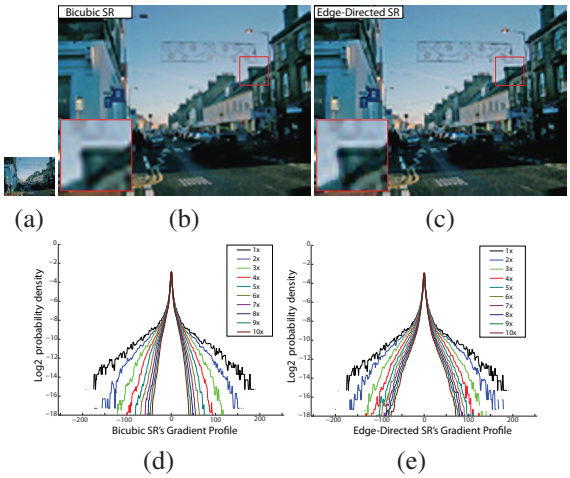

Figure 2. Gradient statistics of HR images using increasing magnification. (a) Input LR image; (b) $10\times$ upsampling using bicubic interpolation; (c) $10\times$ upsampling by edge-directed SR [23]; (d,e) gradient statistics for bicubic interpolation and edge-directed SR with $1\times$ to $10\times$ upsampling. For greater levels of magnification, the gradient statistics increasingly deviate from natural image statistics [8].

namely the gradient profile prior [23]. We then supplement these edges with missing detail taken from a user-supplied example image or texture. The user-supplied texture represents the look-and-feel that the user expects the final super-resolution result to exhibit. To incorporate this detail in a manner consistent with the input image, we also identify significant edges in the example image using the gradient profile prior, and perform a constrained detail transfer that is guided by the edges in the input and example images.

While similar ideas have been used for single image detail- and style- transfer (e.g. [12, 7, 21]), our approach is unique in that it is framed together with edge-directed SR. This gives the user flexibility in specifying the exemplar image – we can still obtain quality edges in the upsampled image even if they are not present in the example image. Experimentally, our procedure produces compelling SR results that are more natural in appearance than edge-directed SR and are on par or better than learning based approaches that require a large database of images to produce quality edges. This is exemplified by the images in Figure 1. The remainder of our paper is organized as follows: Section 2 presents an overview of the reconstruction based formulation; Section 3 discusses details on how edge-directed SR and texture transfer are performed; results are presented in Section 4 with comparisons to previous approaches; Section 5 and 6 conclude our paper with a discussion and summary of our work.

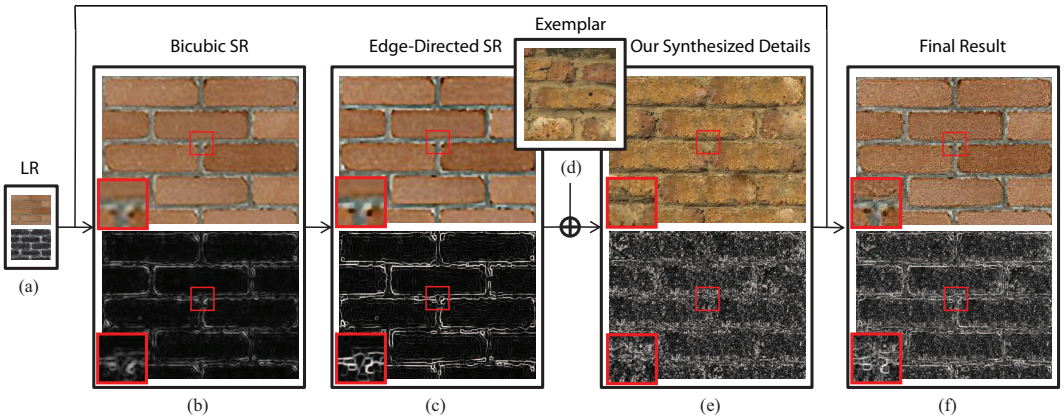

Figure 3. The processing pipeline of our algorithm. (a) Input LR input image with its corresponding gradient profile. (b) Upsampled image and gradient profile using bi-cubic interpolation. (c) Transformed gradient field of (b) using the gradient profile prior [23] to produce sharp SR gradients. (d) Example texture. (e) High resolution gradient field constructed from the high frequency details in (d) with the image structure in (c). (f) Combined gradient field of (c) and (e) used in a reconstruction-based SR to produce the final result.

2. Reconstruction Framework

The processing pipeline of our approach is shown in Figure 3. Given an LR image (Figure 3(a)) and a user supplied image/texture (Figure 3(d)), our goal is to produce a high resolution image (Figure 3(f)) such that its high frequency details resemble those in the example image/texture while preserving the edge structure from the original low resolution image.

Our approach is framed in the standard back-projection formulation typical of reconstruction algorithms [10, 26, 17, 2, 24, 23]. The difference among these various approaches is the prior imposed on the HR image. Our approach is fashioned similar to the gradient profile prior in [23] in which a guidance gradient field, $\nabla_p I_H$, is imposed on the estimated HR image. Unique to our approach is how this $\nabla_p I_H$ is computed. This will be discussed in Section 3.2. First, we describe the main reconstruction algorithm which is necessary for implementation.

Within the reconstruction framework, the goal is to estimate a new HR image, I_H, given the low resolution input image I_L and a target gradient field $\nabla_p I_H$. This can be formulated as a Maximum Likelihood (ML) problem as follows:

$$\begin{aligned} I_H^* &= \arg\max_{I_H} P(I_H|I_L, \nabla_p I_H) \\ &= \arg\min_{I_H} L(I_L|I_H) + L(\nabla_p I_H|\nabla I_H) \\ &= \arg\min_{I_H} ||I_L - d(I_H \otimes h)||^2 + \beta||\nabla_p I_H - \nabla I_H||^2 \end{aligned} \quad (1)$$

where, $L = -logP(\cdot)$, $||I_L - d(I_H \otimes h)||^2$ is the data-cost from the LR image and provides the back-projection constraint, $d(\cdot)$ is the downsampling operator, and $\otimes$ represents convolution with filter h. The term $||\nabla_p I_H - \nabla I_H||^2$ is the data-cost from the guidance gradient field $\nabla_p I_H$, and β is a weight for balancing the two data-costs. Assuming that these data-costs follow a Gaussian distribution, this objective can be cast as a least squares minimization problem with an optimal solution I_H^* obtained by gradient descent with the following iterative update rule [13, 23]:

$$I_H^{t+1} = I_H^t + \tau(I_L - u(d(I_H^t \otimes h)) \otimes p + \beta(\nabla_p^2 I_H - \nabla^2 I_H)) \quad (2)$$

where t is an iteration index, $\otimes$, h, $d(\cdot)$ are defined as in Equation 1, p is the back-projection filter, $u(\cdot)$ is the upsampling process, ∇^2 is the second derivative Laplacian operator and τ is the step size for gradient descent. In the absence of a prior, h and p are chosen to be Gaussian filters with a size proportionate to the super-resolution factor. Satisfactory results are obtained within 30 iterations with $\tau = 0.2$. The parameter β balances the amount of detail in the HR image and the back-projection constraint. The effect of β is demonstrated in Figure 4.

3. Gradient Field Estimation ($\nabla_p I_H$)

The core of our approach involves the transfer of details from the example texture to $\nabla_p I_H$ with respect to structure edges present in I_L. Our approach first upsamples edges from I_L using a reconstruction-based image SR [23]. This is described briefly in Section 3.1 as necessary for implementation; further details can be found in [23]. This edge-directed SR generates sharp edges in the high-resolution target gradient field, and serves as the starting point for our detail synthesis. We will also use this edge-directed SR to identify *structure edges* in the texture example. Details on the constrained texture transfer are provided in Section 3.2.

3.1. Edge-Directed SR via Gradient Profile Prior

Work in [23] has shown that the 1D profile of edge gradients in natural images follows a distribution that is independent of resolution. This so-called gradient profile prior (GPP) provides an effective constraint for upsampling LR images.

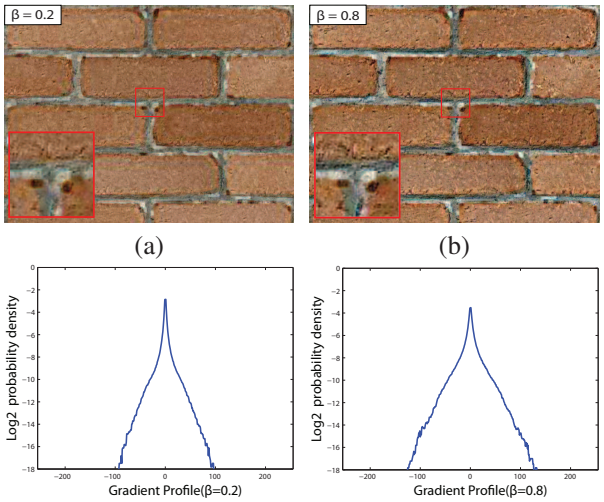

Figure 4. The effect of β on detail synthesis. (a) Results with $\beta = 0.2$; (b) Results with $\beta = 0.8$. To evaluate the amount of detail that has been transferred, we plot the gradient statistics of (a) and (b) in (c) and (d) respectively. The value of β has a direct relationship with the amount of transferred detail.

The gradient profile distribution is modeled by a generalized Gaussian distribution (GGD) as follows:

$$g(x;\sigma,\lambda) = \frac{\lambda\alpha(\lambda)}{2\sigma\Gamma(\frac{1}{\lambda})}\exp(-(\alpha(\lambda)|\frac{x}{\sigma}|)^{\lambda}) \qquad (3)$$

where $\Gamma(\cdot)$ is the gamma function and $\alpha(\lambda) = \sqrt{\Gamma(\frac{3}{\lambda})/\Gamma(\frac{1}{\lambda})}$ is a scaling factor that makes the second moment of the GGD equal to σ^2 and thus allows estimation of σ from the second moment. The parameter λ controls the shape of the generalized Gaussian distribution. Based on a database of over 1000 images, [23] found that the gradient profile distribution of natural images has a shape approximated by a GGD with $\lambda = 1.6$.

To estimate a sharp SR gradient field based on the GPP, we can transform the gradient field of the bicubic upsampled LR image by multiplying the ratio between the gradient profiles of natural images and the gradient profiles of bicubic upsampled LR images as follows:

$$\nabla_g I_H = \frac{g(d;\sigma_h,\lambda_h)}{g(d;\sigma_l,\lambda_l)}\nabla I_L \qquad (4)$$

where $\nabla_g I_H$ is the transformed gradient field, ∇I_L is the gradient field of the bicubic upsampled LR image, d denotes distance of a pixel to an edge maximum, and $g(d;\sigma_h,\lambda_h)$ and $g(d;\sigma_l,\lambda_l)$ represent the learned gradient profiles of natural images and bicubic upsampled images, respectively. After gradient transformation, a sharper and thinner gradient field is obtained as shown in Figure 3(c). This procedure serves as the starting point of our detail synthesis described in the following section.

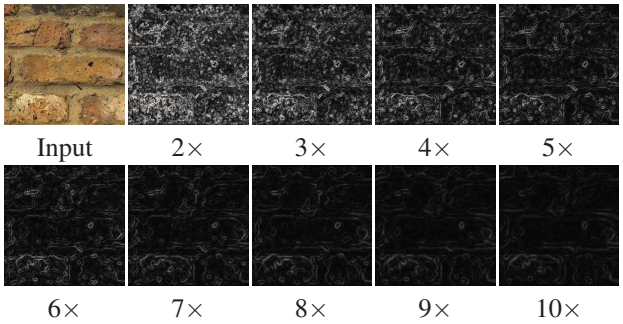

Figure 5. The amount of structure edges $\nabla_g I_E$ versus magnification factor. As the magnification factor increases, the constraints for detail synthesis decrease quadratically, which allows more (larger) details to be transferred to the super-resolution result.

3.2. Synthesis of Details via Example

Given the edge-directed SR gradient field $\nabla_g I_H$ obtained using GPP, and an example image I_E, we now compute the full gradient field prior $\nabla_p I_H$ that includes synthesis of details. By synthesizing details in the gradient domain, issues with illumination and color differences between the LR image and example image are avoided. The input example image I_E represents the look-and-feel for the desired HR image and is assumed to be at the resolution of the HR image. From I_E, example patches are extracted for detail synthesis.

Extracting Structural and Detail Patches In order to better represent edge structure, we extract structure patches from the example image I_E in the following manner. We first downsample I_E to match the scale of the LR image, and then upsample its gradient field using GPP to obtain $\nabla_g I_E$, which represents the salient edge structure in I_E. Note that the amount of extracted structure edges decreases as the magnification factor increases as shown in Figure 5. We now form a set of exemplar patch pairs $\{\nabla E_i, \nabla_g E_i\}$, where *texture patches*, ∇E_i, come directly from ∇I_E and the corresponding *structural patches*, $\nabla_g E_i$, come from the $\nabla_g I_E$. Structural patches $\nabla_g E_i$ are different from ∇E_i, especially as magnification increases.

Detail Synthesis Our detail synthesis is formulated as a constrained texture synthesis using a Markov Random Field (MRF):

$$\begin{aligned}\nabla E^* &= \arg\min_E \sum_i P(\nabla_g I_H|\nabla_g E_i) + \sum_{(i,j)} P(\nabla E_i, \nabla E_j) \\ &= \arg\min_E \sum_i \sum_x ||\nabla_g I_H(x) - \nabla_g E_i(x)||^2 \\ &\quad + \sum_{(i,j)} \sum_{x'\in\Theta} ||\nabla E_i(x') - \nabla E_j(x')||^2 \qquad (5)\end{aligned}$$

where $P(\nabla_g I_H|\nabla_g E_i) = \sum_x ||\nabla_g I_H(x) - \nabla_g E_i(x)||^2$ is the data-cost for aligning structural edges in $\nabla_g E_i$ with the GPP $\nabla_g I_H$, $P(\nabla E_i, \nabla E_j) = \sum_{x'\in\Theta} ||\nabla E_i(x') -$

$\nabla E_j(x')||^2$ is the pairwise energy term to ensure neighborhood patches have similar contents among overlapping regions Θ, $\{x, x'\}$ are local patch coordinates and $\{i, j\}$ are index of nodes in the MRF network.

Since a huge number of exemplar patches can be generated from example image I_E, it is impossible to assign a discrete label to each patch in the MRF process. Therefore, for each image patch location i, we first find the best $K = 15$ candidate exemplar patch pairs that minimize the data term (using the structural patch) and the smoothness term (using the corresponding texture patch). We use patches of size 11×11 that are placed at 7-pixel intervals, providing a 4 pixel overlap. The MRF energy can be optimized using Belief Propagation (BP) [9, 24]. The final result is constructed from the exemplar texture patches, ∇E_i. Structural patches $\nabla_g E_i$ serve to help facilitate better edge alignment in the synthesis process. Feathering is used to blend patches in Θ in the final output of ∇E^*. This optimization procedure for computing ∇E^* is iterated three times, and at each iteration the best $K = 15$ candidate exemplars at each image patch location will be re-evaluated.

Final $\nabla_p I_H$ The final gradient field $\nabla_p I_H$ is then obtained by combining $\nabla_g I_H$ (edge-directed gradient) and ∇E^* (synthesized gradient) as follows:

$$\nabla_p I_H = \begin{cases} \nabla E^*, & \text{if } \nabla E^* \geq \alpha \nabla_g I_H \\ \nabla_g I_H, & \text{otherwise} \end{cases} \quad (6)$$

where α is set to the reciprocal of the magnification factor to maximize detail synthesis. The attenuation factor α is used to counter balance the gradient strengthening effect that edge-directed SR has on $\nabla_g I_H$.

If the user supplies stochastic texture examples with no salient edge structure, the data-cost term will have little effect and the smoothness term will dominate the MRF, resulting in standard texture synthesis. The user may choose to limit the detail synthesis only to selected regions in an image. To facilitate region selection, we currently use a fast interactive image segmentation algorithm [15].

With the estimated $\nabla_p I_H$, we can apply the reconstruction formulation with back-projection as discussed in Section 2 to produce the final HR image.

4. Results

We show results of our algorithm on a variety of examples. In addition, comparisons against other SR approaches are also presented. For all examples, the balance factor in Equation 2 is set as $\beta = 0.5$.

In Figure 1, we compare our approach with GPP [23] for $3\times$ magnification of a monarch butterfly image. An example image was found using Google image search with the query term "Monarch Wing". This example also shows the ground truth image in Figure 1(c). In addition, we show a large $10\times$ magnification in Figure 1(e). Such large magnification appears especially unnatural with edge-directed SR.

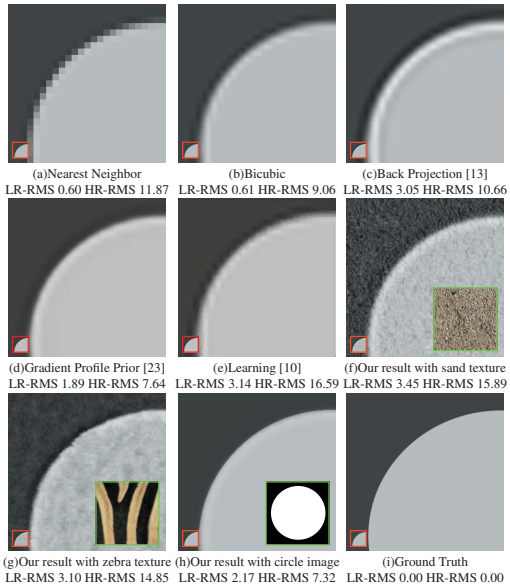

Figure 6. $10\times$ super-resolution on a synthetic example. Our approach generates different results depending on the supplied texture. The lower left corner shows the result image after $10\times$ downsampling. Note that for all results, the down-sampled images are approximately identical. Listed below each result are the LR-RMS errors (RMS errors with respect to the low resolution input), and the HR-RMS errors (RMS errors with respect to the high resolution ground truth image).

Figure 6 shows results with a synthetically generated circle. In this example, the root mean squared (RMS) errors are reported with respect to both the HR image and LR image. For comparison, results with bicubic interpolation, back-projection [13], GPP [23] and Learning [10] are shown. Figure 6(f,g,h) show three examples where different example textures/images have been used. The results exhibit the desired output with details that match the supplied examples. Our method's use of edge-directed SR and constrained detail synthesis produces detail while still preserving edge structure as evident in Figure 6(g,h). Although the results in Figure 6(f) are highly textured, the LR-RMS errors remain small under back-projection. Note that for the Learning approach [10], a generic database is used for super-resolution and hence details in regions are not synthesized. Also, since [10] does not reconstruct high resolution edges before patch matching, some aliasing artifacts remain especially under large scale magnification. This is because using low resolution edges for patch matching contains greater ambiguity. In contrast, our approach uses high resolution edges from reconstruction-based techniques to guide the patch matching, which provides a better and stronger constraint to remove aliasing artifacts. When an example similar to the ground-truth image is used as an example, our method produces a sharper and clearer result

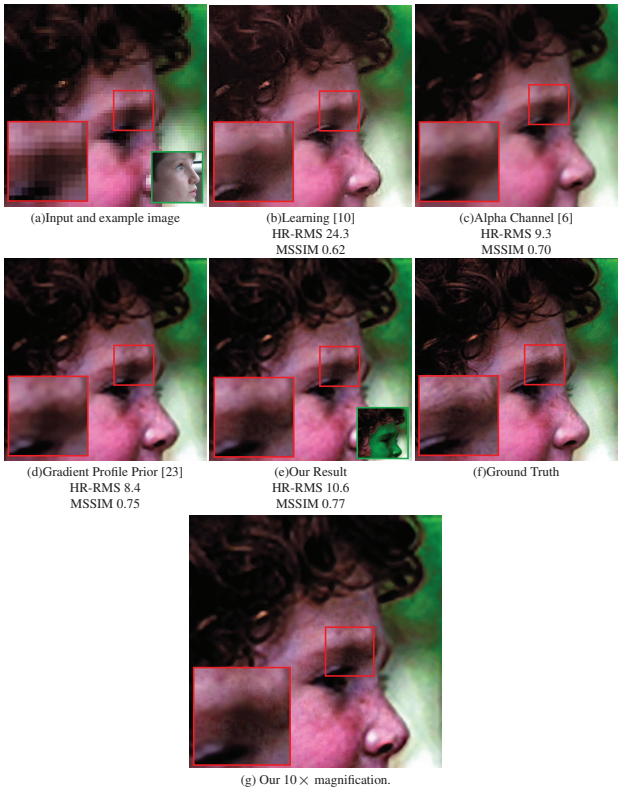

Figure 7. Face with freckles. (a-e) 4× magnification result of various approaches. (f) Ground truth. (g) Our result with a 10× magnification. The HR-RMS errors and the MSSIM score with respect to the 4× ground truth image are listed below each result.

(both subjectively and in terms of RMS errors) as shown in Figure 6(h).

Figure 7 demonstrates SR results for an LR image of a boy's face with noticeable freckles (an example first used in [10]). This image is upsampled with 4× magnification in this experiment. We compare our method against generic learning based SR [10] and two edge-directed techniques ([6] and [23]). Here, we used an image of a different face with significantly different freckle pigmentation to serve as the image example (Google image search "freckle boy" for extra-large images). Our result is shown in Figure 7(e), and a 10× magnification is shown in Figure 7(g). We also compare our result in terms of HR-RMS errors against previous methods. Although our result has larger HR-RMS errors compared with edge-directed techniques ([6] and [23]), our result has much smaller HR-RMS errors compared with generic learning based SR [10]. To better evaluate our result, we show the Mean of Structural Similarity (MSSIM) scores for our results. The MSSIM score is an image quality assessment method that closely matches the human visual system by using local means and variances for measurement [30]. Our result produces the best MSSIM score, because the synthesized details in our result match the "missing" details of the original image in terms of local variances. Previous methods over-smooth the results resulting in lower MSSIM scores.

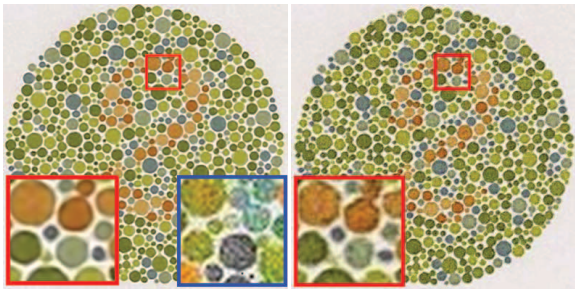

Figure 8. (a) Single image super resolution result from [11] with 3x magnification. The image patch in the blue border is exemplar texture, and the region in the red border is a zoom-in region. (b) Our result which synthesizes details from exemplar texture.

Figure 8 shows a comparison of our result to the single image super resolution approach presented in [11]. From the zoom insets, we can see that the single image approach can produce very nice edges similar to edge-directed approaches (without explicit edge priors). Our result, however, can help synthesize the missing detail to make the result appear more realistic.

Several results under 8× magnification are shown in Figure 9. The LR input image (upsampled using nearest neighbor) and the user-supplied example image are shown in Figure 9(a). Comparisons with GPP [23] (Figure 9(b)) and a standard learning-based approach [10] (Figure 9(c)) are given. Figure 9(d) displays our result with the detail-transfer region shown in the inset and highlighted in green. Each result shows the same zoomed in region for comparison. The images used for example textures were found with Google image search as follows: (first row) "marble texture", (second row) "bark", (third row) "tree sparrow". Our results have sharp edges as well as detail not obtainable with edge-directed SR or standard learning-based SR. Note for the results using [10] we include our example image into the image training database. Even with our example image included with [10], our method still produces better results.

5. Discussion

We have demonstrated that combining edge-directed and learning-based SR obtains good results both subjectively and quantitatively. Our approach produces comparable HR-RMS errors with edge-directed SR but with greater synthesis of details. Since learning-based techniques do not enforce reconstruction constraints, they can have large HR-RMS errors when the synthesized details are not consistent with the LR input and as a result depend heavily on the quality of the training images. Our approach, however, combines learning with guidance from the edge-directed SR

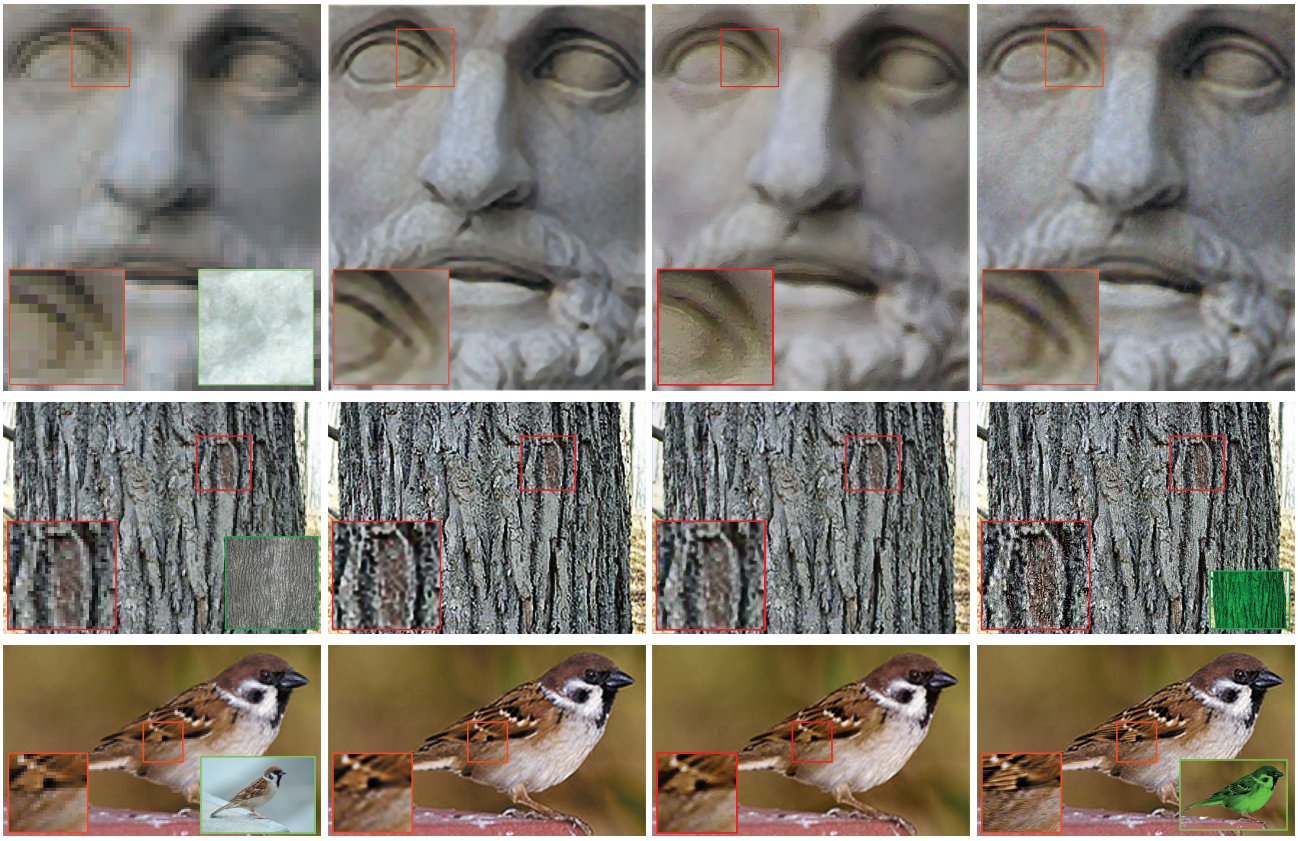

(a) Input (b) Gradient Profile Prior (c) Learning (d) Our Results

Figure 9. Examples with 8× magnification. (a) Input LR image (shown with nearest neighbor upsampling) and an example image/texture provided by the user; (b) results from GPP [23]; (c) results from Learning [10] with a generic database; (d) our results which synthesize details from the example image in the inset of (a). The lower right inset image in (d) highlights regions where details are transferred.

edges to align our synthesized details with the actual structural edges as described in Section 3.2. As a result, we produce SR results that have much smaller HR-RMS errors than purely learning-based techniques. In terms of quality, our results produce more details than pure edge-directed SR, especially for large magnification factors.

The output of our approach is highly dependent on the supplied example image. Since the example serves as a prior for image details, our approach produces very different results with different example images, as illustrated in Figure 6. Although a database of example images could alternatively serve as the prior, a single image example may in fact be easier for the user to control and understand.

This leads to the most obvious question of how to find suitable examples. As demonstrated with our results, suitable examples can generally be found with generic keyword searches using common online image search engines (e.g., Google image search, Yahoo! image search, and photo albums such as Picasa and Flickr). We found that queries with more specific keywords, such as "monarch wing", often yield useful sets of example textures, while more general terms such as "monarch" return numerous (e.g., 100+) images of the desired object in different poses, colors, and resolutions. For certain types of images, finding suitable high-resolution examples may still be challenging, but the variety of images on the web is rapidly growing.

We believe this work suggests a future direction for SR research. Rather than building a generic database from which details are synthesized, we can instead exploit edge-directed SR together with detail transfer to produce results comparable to previous learning-based approaches, even when the supplied example image does not contain suitable edge structure. This strategy is inline with current trends that exploit the availability of rich Internet image data to aid the user in providing examples for the SR problem.

6. Conclusion

We have presented a new framework for image SR that combines edge-directed SR with detail synthesis from a user supplied example image. Our approach uses edge-directed SR to obtain sharp edges by upsampling the LR image, as well as to extract texture structure from the user supplied example. From the example, detail synthesis in the gradient domain is then applied using the edge-directed

HR image. Consistency of the synthesis detail to the input image is then enforced in a reconstruction framework to produce compelling HR images that appear more natural than using learning based or edge-directed SR alone. In addition, our approach is particularly well-suited to leverage the vast example images made available by Internet image search engines and other online image repositories.

7. Acknowledgement

This work was supported in part by a KAIST seed grant (Project No. G04090064) and an NUS AcRF Tier-1 award (Project No. R252000333133).

References

[1] J. Allebach and P. Wong. Edge-directed interpolation. In *ICIP*, 1996.

[2] S. Baker and T. Kanade. Limits on super-resolution and how to break them. *IEEE TPAMI*, 24(9):1167–1183, 2002.

[3] M. Ben-Ezra, Z. Lin, and B. Wilburn. Penrose pixels: Super-resolution in the detector layout domain. In *ICCV*, 2007.

[4] V. Caselles, J. Morel, and C. Sbert. An axiomatic approach to image interpolation. *IEEE TIP*, 7(3):376–386, March 1998.

[5] H. Chang, D. Yeung, and Y. Xiong. Super-resolution through neighbor embedding. In *CVPR*, 2004.

[6] S. Dai, M. Han, W. Xu, Y. Wu, and Y. Gong. Soft edge smoothness prior for alpha channel super resolution. In *CVPR*, 2007.

[7] A. A. Efros and W. T. Freeman. Image quilting for texture synthesis and transfer. In *Proc. ACM SIGGRAPH*, 2001.

[8] R. Fergus, B. Singh, A. Hertzmann, S. T. Roweis, and W. T. Freeman. Removing camera shake from a single photograph. *ACM Trans. Graphics*, 25(3), 2006.

[9] W. Freeman, T. Jones, and E. Pasztor. Example-based super-resolution. *IEEE Computer Graphics and Applications*, 22(2):56–65, 2002.

[10] W. T. Freeman, E. C. Pasztor, and O. T. Carmichael. Learning low-level vision. *IJCV*, 40:25 – 47, 2000.

[11] D. Glasner, S. Bagon, and M. Irani. Super-resolution from a single image. In *ICCV*, 2009.

[12] A. Hertzmann, C. E. Jacobs, N. Oliver, B. Curless, and D. H. Salesin. Image analogies. In *Proc. ACM SIGGRAPH*, 2001.

[13] M. Irani and S. Peleg. Motion analysis for image enhancement: Resolution, occlusion, and transparency. *JVCIR*, 4:324–335, 1993.

[14] X. Li and M. Orchard. New edge-directed interpolation. In *ICIP*, 2000.

[15] Y. Li, J. Sun, C.-K. Tang, and H.-Y. Shum. Lazy snapping. *ACM Trans. Graphics*, 23(3):303–308, 2004.

[16] Z. Lin, J. He, X. Tang, and C.-K. Tang. Limits of learning-based superresolution algorithms. In *ICCV*, 2007.

[17] Z. Lin and H. Shum. Fundamental limits of reconstruction-based superresolution algorithms under local translation. *IEEE TPAMI*, 26(1):83–97, January 2004.

[18] C. Liu, H. Shum, and W. Freeman. Face hallucination: Theory and practice. *IJCV*, 75:115–134, 2007.

[19] C. Liu, H. Y. Shum, and C. S. Zhang. Two-step approach to hallucinating faces: global parametric model and local nonparametric model. In *CVPR*, 2001.

[20] B. Morse and D. Schwartzwald. Image magnification using level-set reconstruction. In *CVPR*, 2001.

[21] G. Ramanarayanan and M. K. Bala. Constrained texture synthesis via energy minimization. *IEEE Transactions on Visualization and Computer Graphics*, 13(1):167–178, 2007.

[22] Q. Shan, Z. Li, J. Jia, and C.-K. Tang. Fast image/video upsampling. *ACM Trans. Graphics*, 27(5), 2008.

[23] J. Sun, J. Sun, Z. Xu, and H. Shum. Image super-resolution using gradient profile prior. In *CVPR*, 2008.

[24] J. Sun, N. N. Zheng, H. Tao, and H. Y. Shum. Generic image hallucination with primal sketch prior. In *CVPR*, 2003.

[25] Y. W. Tai, W. S. Tong, and C. K. Tang. Perceptually-inspired and edge-directed color image super-resolution. In *CVPR*, 2006.

[26] M. F. Tappen, B. C. Russell, and W. T. Freeman. Exploiting the sparse derivative prior for super-resolution and image demosaicing. In *Third International Workshop on Statistical and Computational Theories of Vision*, 2003.

[27] P. Thevenaz, T. Blu, and M. Unser. *Image Interpolation and Resampling*. Academic Press, USA, 2000.

[28] J. van Ouwerkerk. Image super-resolution survey. *Image and Vision Comput.*, 24(10):1039–1052, 2006.

[29] Q. Wang, X. Tang, and H. Y. Shum. Patch based blind image super resolution. In *CVPR*, 2005.

[30] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE TIP*, 13:600 – 612, 2004.