

Image Super-Resolution using Gradient Profile Prior

Jian Sun¹

Jian Sun²

Zongben Xu¹

Heung-Yeung Shum²

¹Xi'an Jiaotong University
Xi'an, P. R. China

²Microsoft Research Asia
Beijing, P. R. China

Abstract

In this paper, we propose an image super-resolution approach using a novel generic image prior – gradient profile prior, which is a parametric prior describing the shape and the sharpness of the image gradients. Using the gradient profile prior learned from a large number of natural images, we can provide a constraint on image gradients when we estimate a hi-resolution image from a low-resolution image. With this simple but very effective prior, we are able to produce state-of-the-art results. The reconstructed hi-resolution image is sharp while has rare ringing or jaggy artifacts.

1. Introduction

The goal of single image super-resolution is to estimate a hi-resolution (HR) image from a low-resolution (LR) input. There are mainly three categories of approach for this problem: interpolation based methods, reconstruction based methods, and learning based methods. The interpolation based methods [12, 29, 18] are simple but tend to blur the high frequency details. The reconstruction based methods [14, 2, 19, 3] enforce a reconstruction constraint which requires that the smoothed and down-sampled version of the HR image should be close to the LR image. The learning based methods [10, 9, 26, 5, 28, 2, 7, 20, 31] “hallucinate” high frequency details from a training set of HR/LR image pairs. The learning based approach highly relies on the similarity between the training set and the test set. It is still unclear how many training examples are sufficient for the generic images.

To design a good image super-resolution algorithm, the essential issue is how to apply a good prior or constraint on the HR image because of the ill-posedness of the image super-resolution. Generic smoothness prior [25, 11] and edge smoothness prior [21, 1, 6, 7, 22, 27] are two widely used priors.

In this paper, we propose a novel generic image prior — *gradient profile prior* for the gradient field of the natural im-

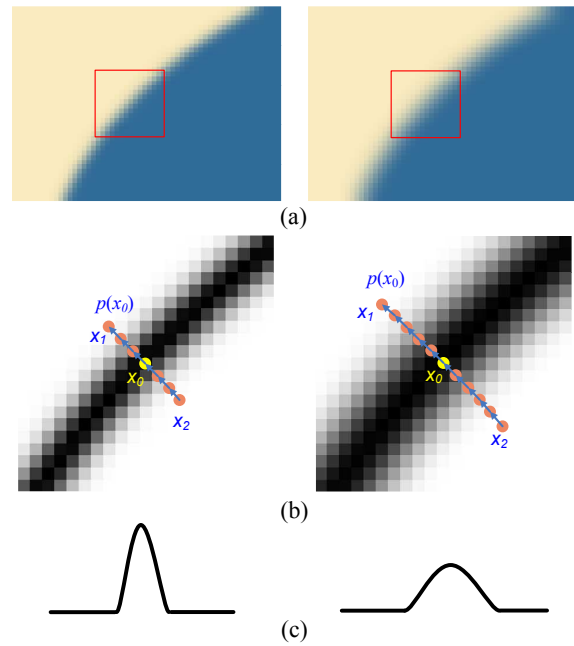


Figure 1. Gradient profile. (a) two edges with different sharpness. (b) gradient maps (normalized and inverted magnitude) of two rectangular regions in (a). $p(x_0)$ is a gradient profile passing through the edge pixel (zero crossing pixel) x_0 , by tracing along gradient directions (two sides) pixel by pixel until the gradient magnitude does not decrease at x_1 and x_2 . (c) 1D curves of two gradient profiles.

age. The gradient profile is a 1-D profile along the gradient direction of the zero-crossing pixel in the image. The gradient profile prior is a parametric distribution describing the shape and the sharpness of the gradient profiles in natural image. One of our observations is that the shape statistics of the gradient profiles in natural image is quit stable and invariant to the image resolution. With this stable statistics, we can learn the statistical relationship of the sharpness of the gradient profile between the HR image and the LR image. Using the learned gradient profile prior and relationship, we are able to provide a constraint on the gradient field of the HR image. Combining with the reconstruction constraint, we can recover a hi-quality HR image.

The advantages of the gradient profile prior are as follows: 1) unlike previous generic smoothness prior and edge smoothness prior, the gradient profile prior is not a smoothness constraint. Therefore, both small scale and large scale details can be well recovered in the HR image; 2) the common artifacts in super-resolution, such as ringing artifacts, can be avoided by working in the gradient domain.

Our work is motivated by recent progresses on natural image statistics. The gradient magnitudes generally obey a heavy tailed distribution e.g., a Laplacian distribution[13]. This kind of “sparseness prior” has been successfully applied to super-resolution[28], denoising [23] [24], inpainting [17], transparency separation [16] and deblurring [8, 15]. However, the sparseness prior only considers the marginal distribution of image gradients (e.g., intensity difference between two adjacent pixels) over the whole image. In this work, our gradient profile prior considers the distribution of the image gradients along local image structures.

Fattal [7] also proposed an edge statistics for image up-sampling. The proposed statistics is the distribution of local intensity continuity in the HR image conditional on edge features in the LR image. Different from his non-parametric statistics, firstly, the gradient profile prior is a generic, parametric image prior for the gradient field of the natural image; secondly, our prior is stable to the image resolution. It is a good property for image super-resolution.

In section 2, we will introduce the gradient profile prior. Then we apply the gradient profile prior to image super-resolution in section 3. We show experimental results in section 4 and conclude the paper in section 5.

2. Gradient Profile Prior

Previous natural image statistics characterizes the marginal distribution of the image gradients over the whole image. The spatial information is discarded. Instead, we study the image gradients along local image structures and the statistical dependency of the image gradients between the HR image and the LR image.

2.1. Gradient profile and its sharpness

Denote the image gradient as $\nabla I = m \cdot \vec{N}$, where m is the gradient magnitude and \vec{N} is the gradient direction. In the gradient field, we denote the zero crossing pixel which is the local maximum on its gradient direction as *edge pixel*.

Figure 1 (a) are two image blocks containing two edges with different sharpness. Figure 1 (b) are corresponding gradient (magnitude) maps. The pixel x_0 in Figure 1 (b) is a zero crossing or edge pixel. Starting from x_0 , we trace a path along the gradient directions (two-sides) pixel by pixel until the gradient magnitude does not decrease anymore. We call the 1-D path $p(x_0)$ as *gradient profile*. Figure 1 (c) are 1D curves of two gradient profiles.

We measure the sharpness of the gradient profile using the square root of the variance (second moment):

$$\sigma(p(x_0)) = \sqrt{\sum_{x \in p(x_0)} m'(x) d^2(x, x_0)} \quad (1)$$

where $m'(x) = \frac{m(x)}{\sum_{s \in p(x_0)} m(s)}$ and $d(x, x_0)$ is the curve length of the gradient profile between x and x_0 . The sharper image gradient profile, the smaller the variance σ is. We call this variance as the profile sharpness.

Profile sharpness estimation. Individually estimating the sharpness for each gradient profile is not robust due to the noise. To have a better estimation, we apply a global optimization to enforce the consistency of neighboring profiles as follows.

First, we construct a graph on all edge pixels. The graph node is the edge pixel and the graph edge is the connection between two neighboring edge pixels within a pre-defined distance (5 pixels in this paper). The edge weight w_{ij} for each clique of two connected nodes i and j is defined as,

$$w_{ij} = \exp(-\zeta_1 \cdot |\nabla u_i - \nabla u_j|^2 - \zeta_2 \cdot d(i, j)^2), \quad (2)$$

where the first term in the exponent is the gradient similarity, and the second term is Euclidean distance between i and j . For each node i , we individually estimate its sharpness $\hat{\sigma}_i$ using Equation (1).

Then, we minimize the following energy to estimate the sharpness of all edge pixels:

$$E(\{\sigma_i\}) = \sum_i [(\sigma_i - \hat{\sigma}_i)^2 + \gamma \cdot \sum_{j \in N(i)} w_{ij} \cdot (\sigma_i - \sigma_j)^2], \quad (3)$$

where $N(i)$ are neighboring nodes of the node i . This energy can be effectively minimized because it is an Gaussian MRF model, in which $\gamma = 5$, $\zeta_1 = 0.15$, and $\zeta_2 = 0.08$ in our implementation.

2.2. Gradient profile prior

Next, we investigate the regularity of the gradient profiles in natural image. We fit the distribution of the gradient profile by a general exponential family distribution, i.e. Generalized Gaussian Distribution (GGD) [30], which is defined as,

$$g(x; \sigma, \lambda) = \frac{\lambda \alpha(\lambda)}{2\sigma \Gamma(\frac{1}{\lambda})} \exp\{-[\alpha(\lambda) \frac{x}{\sigma}]^\lambda\}, \quad (4)$$

where $\Gamma(\cdot)$ is gamma function and $\alpha(\lambda) = \sqrt{\Gamma(\frac{3}{\lambda})/\Gamma(\frac{1}{\lambda})}$ is the scaling factor which makes the second moment of GGD equal to σ^2 . Therefore, σ can be directly estimated using the second moment of the profile. λ is the shape parameter which controls the overall shape of the distribution. The

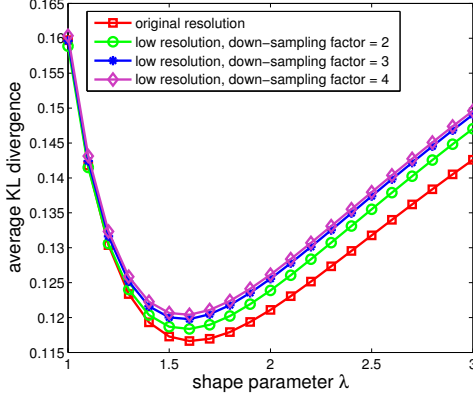


Figure 2. Average KL divergences between the fitted gradient profiles and 1 million gradient profiles by varying the shape parameter λ . The optimal λ is near 1.6 on four data sets with different resolutions.

distribution $g(x; \sigma, \lambda)$ is a Gaussian distribution if $\lambda = 2$, and a Laplacian distribution if $\lambda = 1$.

To fit the distribution, we collect an image set containing 1,000 natural images downloaded from professional photography forums. All images are in the original resolution without down-sampling or up-sampling. For each image, we randomly select 1,000 gradient profiles to construct a data set Ω_1 which consists of 1 million gradient profiles. We also construct other three profile data sets Ω_2 , Ω_3 and Ω_4 from the down-sampled versions of the original resolution images with the down-sampling factors of 2, 3, and 4.

Using Kullback-Leibler (KL) divergence to measure the fitting error, we estimate the optimal λ^* by

$$\lambda^* = \operatorname{argmin}_{\lambda} \left\{ \sum_{p \in \Omega} \text{KL}(p_g(\cdot; \sigma_p, \lambda)) \right\}, \quad (5)$$

where σ_p is the variance (estimated using Equation (3)) of p , which is one profile in the set Ω .

We compute the average KL divergences on four profile sets Ω_1 , Ω_2 , Ω_3 , and Ω_4 by varying the shape parameter λ , as shown in Figure 2. As we can see, the optimal shape parameter is about 1.6 for all down-sampling factors. The shape parameter λ is stable across different resolutions, which means that the gradient profile distribution is resolution-independent in natural image.

We use Pearson's χ^2 hypothesis-test to measure the goodness of our fitted distributions. The χ^2 hypothesis-test for a gradient profile $p(x_0)$ is defined as

$$\chi^2(p) = \sum_{x \in p(x_0)} \frac{[m(x) - E(x)]^2}{E(x)}, \quad (6)$$

where $E(x) = \frac{g(d(x, x_0))}{\sum_{s \in p(x_0)} g(d(s, x_0))} \cdot \sum_{s \in p(x_0)} m(s)$. For significance level κ and degrees of freedom $n - 1$ (n is the number of pixels in p), if $\chi^2(p) < \chi^2_{(\kappa, n-1)}$, the hypothesis that the gradient profile follows the fitted gradi-

ent profile prior cannot be rejected. For the common significance level $\kappa = 0.01$, the average differences between the values of χ^2 on the gradient profiles and corresponding values of $\chi^2_{(\kappa, n-1)}$ are -2.22, -1.90, -1.50, -1.20 on four data sets $\Omega_1, \Omega_2, \Omega_3, \Omega_4$. All average differences are significantly smaller than zero, which means the gradient profiles in natural image are well fitted by our gradient profile prior.

To verify whether the parameter $\lambda = 1.6$ is independent on our collected data or not, we repeat the above experiments on two different image sources. One is 500 images randomly downloaded from Flickr image site. The other is 500 images from a home photo gallery taken with 4 different digital cameras. Again, the obtained optimal shape parameters are stable and between 1.55 and 1.65, which means the generalized gaussian distribution with $\lambda = 1.6$ is a good generic prior for the natural image and independent on the image resolution. Based on this very nice statistics, we only need to study the relationship of the gradient profile sharpness σ between two different resolutions.

2.3. Relationship of gradient profile sharpness between HR image and LR image

Similar to previous methods [10, 26, 7], we study the relationship of gradient profile sharpness between the up-sampled image I_l^u and the HR image I_h , in order to avoid the shifting problem of the zero-crossing pixels in scale space [32]. In our implementation, the up-sampled image I_l^u is the bicubic interpolation¹ of the LR image I_l .

For each gradient profile in the up-sampled image I_l^u , we extract its corresponding gradient profile in the HR image I_h . Because the edge pixels are not exactly aligned in two images, we find the best matched edge pixels by measuring the distance and direction. For each edge pixel e_l in I_l^u , the best matched edge pixel e_h in I_h is found by:

$$e_h = \operatorname{argmin}_{e \in \mathcal{N}(e_l)} \{ \|e - e_l\| + 2\|\vec{N}(e) - \vec{N}(e_l)\| \} \quad (7)$$

where $\mathcal{N}(e_l)$ is the 5×5 neighbors of e_l in the HR image.

To compute the statistics, we quantize the sharpness σ into a number of bins. The width of bin is 0.1. For all LR gradient profiles whose sharpness value falls in the same bin, we calculate the expectation of sharpness of the corresponding HR gradient profiles. Figure 3 shows three fitted curves of computed expectations for the up-sampling factors of 2, 3, and 4. X-axis is the sharpness of the (up-sampled) LR gradient profile and Y-axis is expected sharpness of the hi-resolution gradient profile.

There are two basic observations from Figure 3: 1) the HR gradient profile is sharper than the LR gradient profile

¹Note that the statistic of shape parameter λ in the up-sampled image may be slightly influenced by the bicubic interpolation. However, we found that the optimal λ value for the up-sampled image is still stable. They are 1.63, 1.68, and 1.69 for the up-sampling factors of 2, 3, and 4 on our data sets.

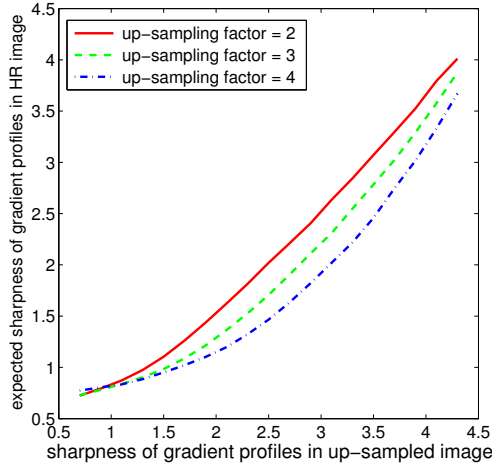


Figure 3. Expected sharpness of the gradient profiles in HR image with respect to sharpness of the corresponding profiles in up-sampled image.

because the bicubic interpolation blurs the profile; 2) the higher the up-sampling factor, the larger the sharpness difference between the HR gradient profile and the LR gradient profile is. Notice that three curves converge together when the sharpness is below 1.0 in Figure 3. One possible reason is due to the inaccuracy of our sharpness estimation. The sharpness estimation for the small scale edge is sensitive to the noise. Also, the introduced image aliasing in the LR image by down-sampling may result in over-estimated sharpness.

3. Gradient Prior for Image Super-Resolution

In this section, we apply the gradient profile prior to image super-resolution. Given a LR image, the gradient profile prior can provide constraints on the gradient field of the HR image: 1) the shape parameter of gradient profiles in the HR image is close to the value 1.6; 2) the sharpness relationship of gradient profiles between two resolutions follows the statistical dependency learned in the previous section. To enforce these constraints, we propose a simple approach as follows.

3.1. Gradient field transformation

We propose a gradient field transformation approach to approximate the HR gradient field by transforming the LR gradient field using the gradient profile prior.

First, we study how to transform a gradient profile $p_l = \{\lambda_l, \sigma_l\}$ in the up-sampled image I_l^u to a gradient profile $p_h = \{\lambda_h, \sigma_h\}$ in the HR image I_h . We compute the ratio

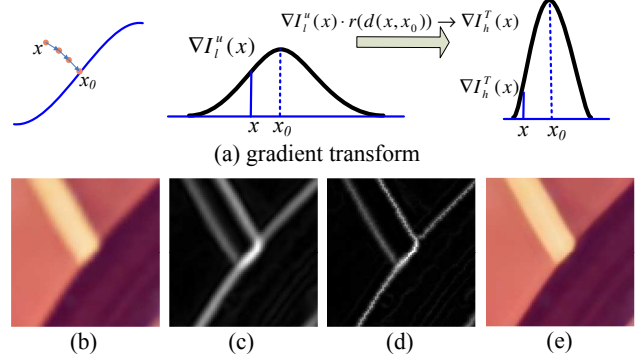


Figure 4. Gradient field transformation. (a) left and middle sub-figures illustrate a gradient profile passing through x and x_0 in the up-sampled image. The gradient of x is transformed to its HR version (right) by multiplying a ratio $r(d(x, x_0))$. (b) and (c) are an up-sampled image and its gradient field. (d) and (e) are transformed gradient field and reconstructed image by solving Poisson equation.

between two gradient profiles, i.e.

$$\begin{aligned} r(d) &= \frac{g(d; \sigma_h, \lambda_h)}{g(d; \sigma_l, \lambda_l)} \\ &= c \cdot \exp\left\{-\left(\frac{\alpha(\lambda_h) \cdot |d|}{\sigma_h}\right)^{\lambda_h} + \left(\frac{\alpha(\lambda_l) \cdot |d|}{\sigma_l}\right)^{\lambda_l}\right\}, \end{aligned} \quad (8)$$

where $c = \frac{\lambda_h \cdot \alpha(\lambda_h) \cdot \sigma_l \cdot \Gamma(1/\lambda_l)}{\lambda_l \cdot \alpha(\lambda_l) \cdot \sigma_h \cdot \Gamma(1/\lambda_h)}$ and d is the curve distance to the edge pixel along the gradient profile. Thus, the HR gradient profile p_h can be estimated by multiplying LR gradient profile p_l by the transform ratio. The shape parameters λ_h and λ_l are set to the learned values in Section 2, the sharpness σ_l is estimated from the image I_l^u and the sharpness σ_h is set as the expected value of σ_l using the relationship we learned in section 2.3.

Second, using the ratio computed in (8), we can transform the LR gradient field ∇I_l^u to the HR gradient field ∇I_h^T by

$$\nabla I_h^T(x) = r(d(x, x_0)) \cdot \nabla I_l^u(x), \quad (9)$$

where x_0 is the edge pixel of the gradient profile passing through x , and $d(x, x_0)$ is the distance between x and x_0 along gradient profile. In our implementation, to find the gradient profile passing through x , we trace from x along the direction (gradient direction or minus gradient direction) with increasing gradient magnitude until reach an edge pixel x_0 (in a threshold distance, e.g., 1 pixel), then adjust the gradient of x by (9).

Figure 4 (a) shows an illustration of gradient transformation. Figure 4 (b-e) gives a real example. Figure 4 (c) is the gradient field of the up-sampled image in Figure 4 (b). Figure 4 (d) is the transformed gradient field and Figure 4 (e) is the reconstructed image by solving Poisson equations. The recovered image is sharp and with rare ringing artifacts.

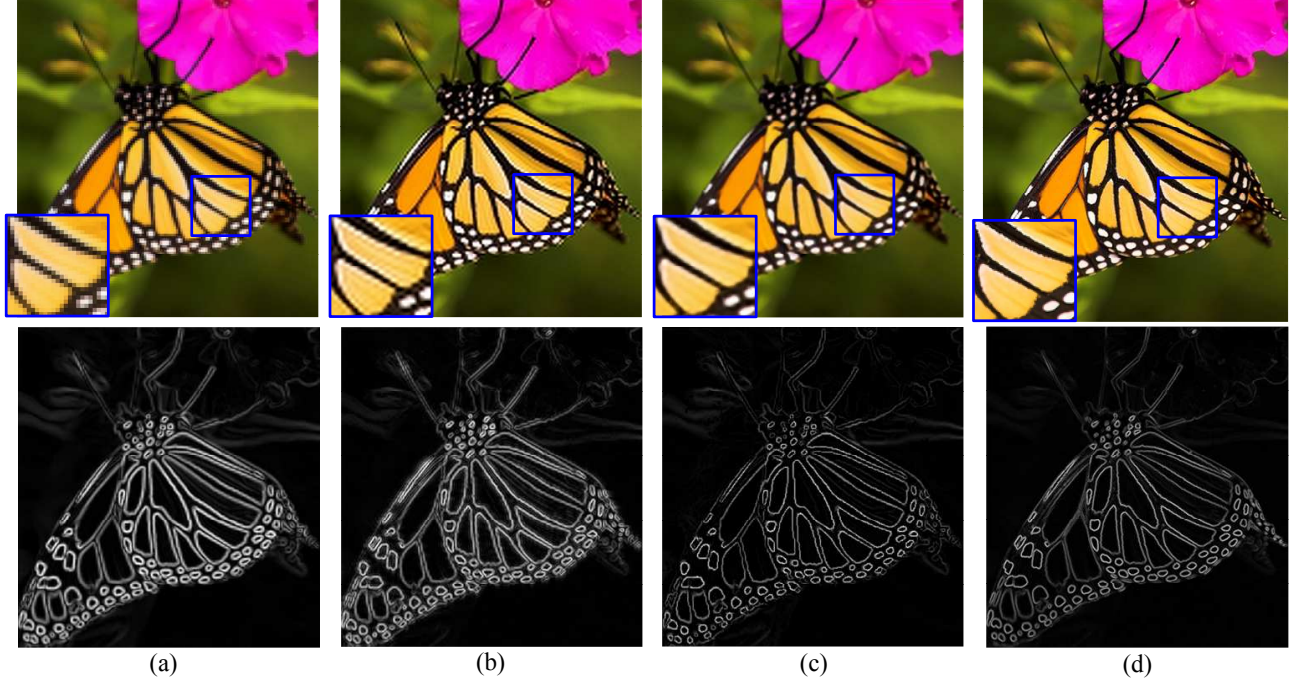


Figure 5. HR image reconstruction (3X). (a) LR image (nearest neighbor interpolation) and gradient field of its up-sampled image (bicubic interpolation). (b) result of back-projection and its gradient field, (c) our result and transformed gradient field for HR image. (d) ground truth image and its gradient field. Compared with the gradient field of result by back-projection, the transformed gradient field is much closer to the ground truth gradient field of HR image. Our reconstructed result has rare jaggy or ringing artifacts.




Figure 6. Super-resolution on synthetic image (4X). (a) LR image (nearest neighbor interpolation). (b) reconstructed HR image. (c) gradient field of the up-sampled image (bicubic interpolation), (d) transformed gradient field from (c).

3.2. HR Image reconstruction

We use the transformed gradient field as the gradient domain constraint for the HR image reconstruction. Given the LR image I_l , in order to reconstruct the HR image I_h , we minimize the following energy function by enforcing the constraints in both image domain and gradient domain:

$$E(I_h|I_l, \nabla I_h^T) = E_i(I_h|I_l) + \beta E_g(\nabla I_h|\nabla I_h^T), \quad (10)$$

where $E_i(I_h|I_l)$ is the reconstruction constraint in the image domain and $E_g(\nabla I_h|\nabla I_h^T)$ is the gradient constraint in the gradient domain.

The reconstruction constraint measures the difference between the LR image I_l and the smoothed and down-

sampling version of HR image I_h , i.e.

$$E_i(I_h|I_l) = |(I_h * G) \downarrow - I_l|^2. \quad (11)$$

where G is a spatial filter, $*$ is the convolution operator, and \downarrow is the down-sampling operation. We use a gaussian filter for the spatial filter G . The kernel standard variance is set to 0.8, 1.2 and 1.6 for the up-sampling factors of 2, 3 and 4.

The gradient constraint requires that the gradient field of the recovered HR image should be close to the transformed HR gradient field ∇I_h^T :

$$E_g(\nabla I_h|F) = |\nabla I_h - \nabla I_h^T|^2, \quad (12)$$

where ∇I_h is the gradient of I_h . Using this constraint, we encourage the gradient profiles in I_h has a desired statistics which we learned from the natural images.

The energy (10) can be minimized by a gradient descent algorithm:

$$I_h^{t+1} = I_h^t - \tau \cdot \frac{\partial E(I_h)}{\partial I_h},$$

where

$$\frac{\partial E(I_h)}{\partial I_h} = ((I_h * G) \downarrow - I_l) \uparrow * G - \beta \cdot (\nabla^2 I_h - \nabla^2 I_h^T). \quad (13)$$




Figure 7. Super-resolution comparison (3X). Gradient reconstruction is obtained by solving poisson equations on the transformed gradient field. Both of gradient reconstruction result (e) and our result (f) contain much less ringing artifacts, especially along the image edges. But our result (f) is closer to the ground truth by enforcing the reconstruction constraint. See text for details.



Figure 8. Super-resolution comparison (4X) of learning based method [10], alpha channel super-resolution [6], and our approach. Both large scale edges and small scale details (on the face) are recovered in our result.

The global optimum can be obtained because the energy (10) is a quadratic function. We set the step size τ to 0.2, parameter $\beta = 0.5$ and use the up-sampled image I_l^u as the initial value of I_h .

Figure 5 gives a real example of our method. Figure 5 (a) are input LR image and the gradient field of bicubic up-sampled image. Figure 5 (d) are ground truth HR image and its gradient field. Figure 5 (b) are back-projection[14] result using the reconstruction constraint only. Notice the ringing artifacts in both image and gradient field. The bottom image in Figure 5 (c) is our transformed gradient field. As we can see, it is much closer to the ground truth gradient field shown in Figure 5 (d). The top image in Figure 5 (c) is our final reconstructed HR image using both image and gradient constraints. The ringing artifacts are substantially suppressed by the gradient constraint.

Figure 6 also shows an example on a synthetic image. Our approach can reconstruct a very sharp HR image guided by a transformed gradient field.

4. Experiments

We test our approach on a variety of images. For the color images, we only perform image super-resolution on the grayscale channel because the human are more sensitive to the brightness information. The color channels are up-sampled using the bicubic interpolation.

In Figure 7, we compare our approach with bicubic interpolation, sharpened bicubic interpolation, back-projection [14], and reconstruction from the transformed gradient field by solving poisson equations. The result of bicubic interpolation is over-smooth, for example the region in the rectangle. The sharpened bicubic interpolation and back-projection introduce ringing or jaggy artifacts, especially along salient edges. The result of reconstruction from the transformed gradient field is sharp and with rare artifacts, but the color is not close to the ground truth HR image. By combining gradient constraint and reconstruction constraint, our final result is the best.

Figure 8 shows the comparison of our approach with learning based method [10] and alpha channel super-resolution [6]. The result of learning based method is sharp in appearance. However, high frequency artifacts are also introduced from the training samples, for example the artifacts around the nose. The salient edges in alpha channel super-resolution result are sharp, but the small scale edges, for example flecks on the face, are not well recovered. That's because it's hard to estimate alpha channel value for the edges with weak contrast and large blur. Compared with these results, our approach can recover both large scale edges and small scale details, and introduce minimal additional artifacts.

Figure 9 and 10 show four examples with up-sampling




Figure 9. Super-resolution results with up-sampling factors of 8 and 16.




Figure 10. More super-resolution results with up-sampling factor of 8. The left image is the LR image, and the right image is our result.

factor of 8 and one example with up-sampling factor of 16, in which the HR results are produced by repeatedly running

our super-resolution algorithm with up-sampling factor of 2. In Figure 9, the image regions in the blue rectangles are magnified by nearest neighbor interpolation for better illustration. All of the results show that our method can reliably recover the image details and produce sharp edges with minimal additional artifacts.

We also compute RMS and ERMS [26] to qualitatively measure the super-resolution results of Monarch (Figure 5), Lena (Figure 7) and Head (Figure 8). The measurements are listed in Table 1. Our model outperforms the bicubic and back-projection with lower RMS and ERMS. The computation costs for Monarch (original resolution is 399×423), Lena (original resolution is 500×500) and Head (original resolution is 280×280) are 7.4s, 8.7s, and 3.5s on a 3.0 GHz PC.

Table 1. Super-resolution quality measurement.

test images	bicubic		back-projection		our method	
	RMS	ERMS	RMS	ERMS	RMS	ERMS
Monarch	16.4	26.0	13.6	21.3	13.2	20.9
Lena	8.8	11.5	8.2	10.8	7.8	10.1
Head	8.7	10.9	8.6	10.6	8.4	10.3

5. Conclusion and Discussion

In this paper, we have established a gradient profile prior for natural image. Using this prior, a gradient field constraint is enforced for the problem of image super-

resolution. The gradient constraint helps to sharpen the details and suppress ringing or jaggy artifacts along edges. Encouraging results are obtained on a variety of natural and synthetic images.



(a) noisy input

(b) our result

Figure 11. Super-resolution on a noisy image (4X). Noisy LR image is denoised by non-local denoising method [4], then the denoised image is up-sampled by the proposed method, and the noises are up-sampled by bilinear interpolation.

For noisy input LR image, estimating the gradient profile might be inaccurate due to the noise. One possible solution is to denoise the LR image first, then add the up-sampled noises back after the image super-resolution, see Figure 11 for an example. In the future, we are planning to extend the proposed method to video super-resolution. We are also interested in applying the gradient profile prior to other image reconstruction applications.

Acknowledgments

We thank the anonymous reviewers and area chairs for helping us to improve this paper. This work is performed when the first author visited Microsoft Research Asia. The first author and Zongben Xu were supported by the National Basic Research program (973 project) of China (No. 2007CB311002).

References

- [1] H. A. Aly and E. Dubois. Image up-sampling using total-variation regularization with a new observation model. *IEEE Trans. on IP*, 14(10):1647–1659, 2005.
- [2] S. Baker and T. Kanade. Limits on super-resolution and how to break them. *IEEE Trans. on PAMI*, 24(9):1167–1183, 2002.
- [3] M. Ben-Ezra, Z. C. Lin, and B. Wilburn. Penrose pixels: Super-resolution in the detector layout domain. In *ICCV*, 2007.
- [4] A. Buades, B. Coll, and J. M. Morel. A non-local algorithm for image denoising. In *CVPR*, volume 2, pages 60–65, 2005.
- [5] H. Chang, D. Y. Yeung, and Y. Xiong. Super-resolution through neighbor embedding. In *CVPR*, volume 1, pages 275–282, 2004.
- [6] S. Y. Dai, M. Han, W. Xu, Y. Wu, and Y. H. Gong. Soft edge smoothness prior for alpha channel super resolution. In *CVPR*, 2007.
- [7] R. Fattal. Image upsampling via imposed edge statistics. *ACM Transactions on Graphics*, 26(3):95:1–95:8, 2007.
- [8] R. Fergus, B. Singh, A. Hertzmann, S. T. Roweis, and W. T. Freeman. Removing camera shake from a single photograph. *ACM Transactions on Graphics*, 25(3):787–794, 2006.
- [9] W. T. Freeman, T. R. Jones, and E. C. Pasztor. Example-based super-resolution. *IEEE Computer Graphics and Applications*, 22(2):56–65, 2002.
- [10] W. T. Freeman, E. Pasztor, and O. Carmichael. Learning low-level vision. *International Journal of Computer Vision*, 40(1):25–47, 2000.
- [11] R. C. Hardie, K. J. Barnard, and E. Armstrong. Joint map registration and high-resolution image estimation using a sequence of undersampled images. *IEEE Trans. on IP*, 6(12):1621–1633, 1997.
- [12] H. S. Hou and H. C. Andrews. Cubic splines for image interpolation and digital filtering. *IEEE Trans. on SP*, 26(6):508–517, 1978.
- [13] J. G. Huang and D. Mumford. Statistics of natural images and models. In *CVPR*, volume 1, pages 541–547, 1999.
- [14] M. Irani and S. Peleg. Motion analysis for image enhancement: Resolution, occlusion and transparency. *Journal of Visual Communication and Image Representation*, 4(4):324–335, 1993.
- [15] A. Levin, R. Fergus, F. Durand, and W. T. Freeman. Image and depth from a conventional camera and depth from a conventional camera with a coded aperture. *ACM transactions on Graphics*, 26(3):70:1 – 70:9, 2007.
- [16] A. Levin, A. Zomet, S. Peleg, and Y. Weiss. Seamless image stitching in the gradient domain. In *ECCV*, volume 4, pages 377 – 389, 2005.
- [17] A. Levin, A. Zomet, and Y. Weiss. Learning how to inpaint from global image statistics. In *ICCV*, volume 1, pages 305– 312, 2003.
- [18] X. Li and M. T. Orchard. New edge-directed interpolation. *IEEE Trans. on IP*, 10(10):1521–1527, 2001.
- [19] Z. C. Lin and H. Y. Shum. Fundamental limits of reconstruction-based superresolution algorithms under local translation. *IEEE Trans. on PAMI*, 26(1):83–97, 2004.
- [20] C. Liu, H. Y. Shum, and W. T. Freeman. Face hallucination: Theory and practice. *International Journal of Computer Vision*, 75(1):115–134, 2007.
- [21] B. S. Morse and D. Schwartzwald. Image magnification using level-set reconstruction. In *CVPR*, volume 1, pages 333–340, 2001.
- [22] V. Rabaud and S. Belongie. Big little icons. In *CVPR*, volume 3, pages 24–30, 2005.
- [23] S. Roth and M. J. Black. Fields of experts: A framework for learning image priors. In *CVPR*, volume 2, pages 860–867, 2005.
- [24] S. Roth and M. J. Black. Steerable random fields. In *ICCV*, 2007.
- [25] R. Schultz and R. Stevenson. Extraction of high-resolution frames from video sequences. *IEEE Trans. IP*, 5(6):996 – 1011, 1996.
- [26] J. Sun, N. N. Zheng, H. Tao, and H. Y. Shum. Image hallucination with primal sketch priors. In *CVPR*, volume 2, pages 729–736, 2003.
- [27] Y. W. Tai, W. S. Tong, and C. K. Tang. Perceptually-inspired and edge-directed color image super-resolution. In *CVPR*, volume 2, pages 1948–1955, 2006.
- [28] M. F. Tappen, B. C. Russell, and W. T. Freeman. Exploiting the sparse derivative prior for super-resolution and image demosaicing. In *IEEE Workshop on Statistical and Computational Theories of Vision*, 2003.
- [29] P. Thevenaz, T. Blu, and M. Unser. *Image Interpolation and Resampling*. Handbook of Medical Imaging, Processing and Analysis, Academic Press, San Diego, USA, 2000.
- [30] M. K. Varanasi and B. Aazhang. Parametric generalized gaussian density estimation. *The Journal of the Acoustical Society of America*, 86(4):1404–1415, 1989.
- [31] Q. Wang, X. Tang, and H. Y. Shum. Patch based blind image super resolution. In *ICCV*, volume 1, pages 709–716, 2005.
- [32] A. L. Yuille and T. Poggio. Fingerprints theorems for zero crossings. *J. Opt. Soc. Am. A*, 2:683–692, 1985.