

Irradiance Preserving Image Interpolation

Andrea Giachetti

VIPS lab, Dipartimento di Informatica, Università di Verona

E-mail: andrea.giachetti@univr.it

Abstract

In this paper we present a new image upscaling (single image superresolution) algorithm. It is based on the refinement of a simple pixel decimation followed by an optimization step maximizing the smoothness of the second order derivatives of the image intensity while keeping the sum of the brightness values of each subdivided pixel (i.e. the estimated irradiance on the area) constant. The method is physically grounded and creates images that appear very sharp and with reduced artifacts. Subjective and objective tests demonstrate the high quality of the results obtained.

1. Introduction

Single image super resolution is a hot topic in the Computer Graphics and Image Processing communities. Upscaling algorithms are, in fact, widely applied, for example to enhance printed images or to display low quality images and videos on high resolution displays. As pointed out by in a recent review [15], there are three main problems with simple kernel-based interpolation: (i) it creates oversmoothed images, (ii) generates jagged artifacts (iii) it is not able to guess reasonable high-frequency components from the original data. While the first problem can be reduced by applying a sharpening filter (or directly using a Lanczos kernel that enhances intensity discontinuities), more complex algorithms should be applied in order to reduce the effects of the other two. Edge-directed methods adapt the local interpolation to the estimated local edge behavior. They are often based on filling schemes putting the original ones in an enlarged grid and filling the holes with weighted averages of the neighboring pixels with weights depending on the edge features [1, 8, 10]. They provide images with reduced jagged artifacts, but often oversmoothed and in some cases affected by other kind of artifacts in high-frequency regions.

Example-based methods [4, 6, 7, 12] try learning the relationship between low-resolution and high-resolution patches from a training set of images. This is accomplished by reconstructing the high-resolution im-

age by merging detailed patches corresponding to the coarse ones. These methods can provide natural and sharp images, and are obviously able to guess reasonable (but not necessarily correct) high-frequency components. Their drawbacks are related to the computational complexity, the necessity of a representative training set and the risk of having high frequency components that do not correspond to the true scene. Learning methods can be applied to adapt interpolation coefficients to the edge behavior. In the resolution synthesis method [2], for example, low resolution pixels are first classified in the context of a window of neighboring pixels. Then the corresponding high-resolution pixels are obtained by filtering with coefficients depending on the classification result.

Another class of methods is based on optimization techniques. Ad hoc constraints are used to define energy functions that should be minimized when the high resolution image is, in some sense the most probable given the low resolution one. In any case, there are several methods that differ mostly in the way they impose edge continuity and sharpness. In [11] a gradient profile prior derived from the analysis of natural images and relating gradient profiles at different scales is used to enhance sharpness. In [9] a constraint related to the smoothness of isophote curves is applied. In [13] the Gaussian Point Spread Function in the classical backprojection scheme is locally modified according to a local multiscale edge analysis. In [5] after the use of a grid filling scheme, the added pixels are refined with the constraints related to the edge curvature continuity, trying at the same time to maximize the gradient components. In [14] the process generating the high resolution image is explicitly modeled as the recapturing of the scene model in a Bayesian inference modeling.

The proposed method adopts an optimization scheme that in some sense simulates the capture of the scene with a different sensor. The idea is to upscale images with an integer factor assuming the constancy of the irradiance incident on the original pixel area. For this reason, the proposed approach is denoted

Irradiance-Preserving Interpolation (IPI). The optimization depends on a second order derivative continuity prior ensuring edge continuity and smoothness.

2. Algorithm description

Upscaling an image of an integer factor F is equivalent to splitting each pixel into $F \times F$ smaller ones. Ideally, the procedure should recover the image that would have been acquired by a different real sensor with pixels of this new smaller size.

We assume that the light incidence over the original pixel and its corresponding set of $F \times F$ smaller new pixels should be the same.

If $I = g(W/S)$ is the formula giving the local gray level as a function of the irradiance incident (W) on the pixel surface s , and if $I(i, j) = g(W(i, j)/s_0)$ is the gray level of the low resolution pixel, the gray levels of the $F \times F$ decimated pixels will be expressed by $I(i_h, j_k) = g(W(i_h, j_k)/s)$ with $s = S/(F \times F)$. The constraint of equal energy can be introduced as follows:

$$\sum_h \sum_k W(i_h, j_k) = F^2 W(i, j) \quad (1)$$

The image upscaling algorithm can then change the single $I(i_h, j_k)$ values, while keeping the sum of the related irradiances constant and equal to the irradiance incident on the original pixel:

$$\sum_h \sum_k W(i_h, j_k) = \sum_h \sum_k g^{-1}(I(i_h, j_k)) = F^2 W(i, j) \quad (2)$$

In the simplified hypothesis that $g^{-1}(I)$ is a simple linear function, we can rewrite the constraint as the requirement that the sum of the new virtual pixel values is exactly the level of the original times the squared scale factor:

$$\sum_{h=1}^F \sum_{k=1}^F I(i_h, j_k) = F^2 I(i, j) \quad (3)$$

Our goal is to introduce an optimization procedure creating a "natural" edge behavior in the image while satisfying this constraint. It has been shown that a refinement of the upscaled image imposing the continuity of the second order derivatives of the image intensity is effective in reducing the jagged artifacts. The Iterative Curvature Based Interpolation (ICBI) algorithm[5], based on this idea, creates artifact-free images, even if oversmoothed.

We therefore decided to add this constraint to (3), in order to create an optimization procedure able to create sharp upscaled images with reduced artifacts. The idea is to apply, after the initial pixel splitting, different perturbations of the values inside each original pixel

that keep sum of the new pixels constant and select then the one minimizing an energy function that is just the sum of the difference between the second order derivatives (horizontal and vertical) inside and outside at the boundaries of the large pixel region. Let us describe in detail the procedure for a $2 \times$ enlargement. For each original pixel we have four new pixels, initialized with the same intensity level (or color components) as the big one. Then, for each 2×2 block $B(i, j)$ of the new image corresponding to an original pixel, we compute 9 possible modified blocks by adding to them the following matrices M_k :

$$\begin{aligned} M_1 &= \alpha \begin{bmatrix} 3 & -1 \\ -1 & -1 \end{bmatrix} & M_2 &= \alpha \begin{bmatrix} -1 & 3 \\ -1 & -1 \end{bmatrix} & M_3 &= \alpha \begin{bmatrix} -1 & 3 \\ -1 & -1 \end{bmatrix} \\ M_4 &= \alpha \begin{bmatrix} 3 & -1 \\ -1 & -1 \end{bmatrix} & M_5 &= \alpha \begin{bmatrix} -1 & 3 \\ -1 & -1 \end{bmatrix} & M_6 &= \alpha \begin{bmatrix} -1 & 3 \\ -1 & -1 \end{bmatrix} \\ M_7 &= \alpha \begin{bmatrix} 3 & -1 \\ -1 & -1 \end{bmatrix} & M_8 &= \alpha \begin{bmatrix} -1 & 3 \\ -1 & -1 \end{bmatrix} & M_9 &= \alpha \begin{bmatrix} -1 & 3 \\ -1 & -1 \end{bmatrix} \end{aligned}$$

At each iteration we compute for each modified configuration $B(i, j)_k$ the local energy term:

$$\begin{aligned} E(i, j) &= \\ & \sum_{x, y \in B_{ij}} \left| \frac{\partial H_{11}(I_k(x, y))}{\partial x} \right| + \sum_{x, y \in B_{ij}} \left| \frac{\partial H_{11}(I_k(x, y))}{\partial y} \right| \\ & \sum_{x, y \in B_{ij}} \left| \frac{\partial H_{22}(I_k(x, y))}{\partial x} \right| + \sum_{x, y \in B_{ij}} \left| \frac{\partial H_{22}(I_k(x, y))}{\partial y} \right| \end{aligned} \quad (4)$$

where H_{11} and H_{22} are the diagonal elements of the the Hessian matrix and I_k is the current high-resolution image, modified only in the block (i, j) by adding M_k . We finally update each block replacing its pixel values with the modified block B_{min} corresponding to the lower energy.

We proceed with this greedy minimization until the brightness change is lower than a threshold (or for a fixed number of iterations if time constraints are present). Different upscaling factors can be obtained in a similar way (with a larger number of larger masks M_k). As in NEDI-like techniques, $2^N \times$ factors can be obtained by applying repeated $2 \times$ zooming, and we used this method to obtain the $4 \times$ upscaling in the experiments described in the next Section.

We also implemented a different greedy optimization scheme, minimizing at all high-resolution pixel locations an energy term depending on the derivatives and on a new term that is a function of the difference between the average of neighboring pixels belonging to the same low-resolution pixel and the low-resolution value. In this way the irradiance preservation constraint is introduced in a soft way, depending on the function choice. After several tests we applied a function that is zero if the difference is lower than a threshold and goes




Figure 1. Basic idea of the IPI method: pixels are first decimated, then the optimization procedure tries to minimize the variation of second-order image derivatives at pixel borders while keeping the sum of the values constant.

rapidly to infinity if higher. Results, however, are not so good as with the “strict” constraint.

3. Experimental results

We analyzed the performances of the method on different kind of images. Large image databases are not widely applied to compare upscaling methods. We tried the method on the 25 images used in [1], available at the web site cited in the paper, using the same comparison code provided with the data (cropping out a 50 pixel border and computing PSNR, using ad hoc shifted reference images for odd super-resolution). In the test images rescaled to 256×256 are upscaled by a factor 2 and compared with reference images. Images rescaled to 128×128 are upscaled by a factor 4 and compared with reference images. We did not limit the comparison to the estimation of Peak Signal To Noise Ratio (PSNR), but also tested other quality measures. As in [15] we evaluated the Edge Stability Measure (ESMSE) introduced in [3], where it is described as a very good way to detect artifacts. We also evaluated artifacts by counting pixels with intensity differing more than a threshold (set to 10). We called it LDPP (Largely Different Pixels Percentage): its value is also related to possible errors of labelling algorithms based on intensity. Table 1 shows the results obtained with the IPI method compared with other kernel-based (Bicubic, BC; Lanczos, LA) and edge-based methods (iNedi, [1]). The IPI implementation with the strict irradiance-preserving constraint always got the best scores. The alternative implementation (IPI2) with a relaxed constraint did not reach the same accuracy, even if we tried several options for the irradiance related constraint. This is probably due to the fact that the simultaneous variation of all the subdivided pixels provides a better greedy minimization less influenced by local minima.

From a visual point of view, the ability of the method to create sharper images than those provided by other methods is clear. In Fig. 2 the quality of the reconstructed detail is evident compared with the result

Peak Signal to Noise Ratio (dB)						
algorithm	BC	LA	iNEDI	ICBI	IPI	IPI2
2X	30.36	31.06	30.64	31.13	32.33	31.51
4X	24.91	25.29	25.17	25.38	26.16	25.67
Edge Stability Measure						
algorithm	BC	LA	iNEDI	ICBI	IPI	IPI2
2X	0.65	0.54	0.81	0.61	0.41	0.50
4X	2.65	2.53	2.72	2.68	2.26	2.47
Largely Different Pixel Percentage(%)						
algorithm	BC	LA	iNEDI	ICBI	IPI	IPI2
2X	8.7	7.8	8.0	7.6	6.2	7.0
4X	18.0	17.2	17.1	17.0	15.2	16.0

Table 1. Average error values obtained comparing 25 test images with their downscaled versions upscaled of the same factor (2x or 4x) with different algorithms. IPI provides always the best scores.

of iNEDI. In Fig. 3 we compare the same detail of the Lena image reduced to 256×256 and upscaled to 512×512 shown in [15]. The results obtained with IPI (right) is clearly sharper and with more realistic detail than the one obtained with Lanczos kernel. For instance, even if we could not get the upscaling and comparison code of the paper, the PSNR obtained is also far better than all those obtained with the techniques compared in that review paper (35.21 for IPI, 34.13 for Resolution Synthesis [2]).

Our new method is not so efficient as other iterative methods like ICBI, but is faster than methods like NEDI or example-based methods. The average computation time for $2 \times$ upscaling of 256×256 images was about 90 seconds with a non-optimized Matlab implementation on a Dell XPS M1210 with an Intel Core2 Duo T7200 2.0 GHz CPU. In our experience this means that a GPU implementation on a recent graphics card could reach real time performances. The great advantages in terms of error measures does not mean that the method is optimal in any respect. We found that IPI is not so good as ICBI, in order to remove jagged artifacts. Fig. 3 shows a detail of the Baboon image downsized to 256×256 and upscaled by a factor two with ICBI and IPI methods. The new technique provides sharper images, but presents visible jaggies.

The Irradiance Constancy constraint can also be reasonably applied to interpolate medical images reducing the classical “partial volume” effect. We tested the method on a set of 11 medical images of different kinds (X-Ray, MRI, CT, US) and we also obtained for $2 \times$ upscaling of subsampled images compared with original images an average PSNR more than 1dB higher than



Figure 2. Detail of an image of the test dataset enlarged by a factor 4. IPI (center) provides estimated high frequency details clearly closer to the real ones (right) than those provided by other edge-based method such as improved NEDI (left).




Figure 3. Visual comparison of an enlarged detail of the Lena image reduced to 256×256 and enlarged by a factor 2 with Nearest Neighbor interpolation (left), Lanczos Kernel (center), IPI (right). The new method provides the sharpest image and reconstructs reasonable high-frequency details.

those obtained with other techniques. We are currently testing the method on 3D data to interpolate slices or to enhance the resolution of volume visualization.

4 Discussion

We presented a simple physically grounded method to obtain super-resolved images which look natural and sharp. The method is based on an optimization scheme which tries to smooth image derivatives while keeping the irradiance on the original pixels area constant. The results are clearly sharper and more naturally detailed than those obtained with other state-of-the-art kernel-based or edge-based methods, and the quality measures applied in previous works show that the method is a big improvement over previous ones. The method could, however, still be improved in order to have a better removal of jaggies. Another thing we plan to do is to introduce new quality measures able to show a better correlation with visual artifact removal.

Acknowledgements: A. Brughitta checked the English.

References

- [1] N. Asuni and A. Giachetti. Accuracy improvements and artifacts removal in edge based image interpolation. In




Figure 4. Visual comparison of an enlarged detail of the Baboon image reduced to 256×256 enlarged by a factor 2 with nearest neighbor interpolation (left), ICBI (center) and IPI (right). The new method provides the sharpest image, but is not as effective as ICBI in removing jaggies and creating continuous lines.

Proc. 3rd Int. Conf. Computer Vision Theory and Applications (VISAPP'08), 2008.

- [2] C. B. Atkins, C. A. Bouman, and J. P. Allebach. Optimal image scaling using pixel classification. In *Proc. IEEE Int. Conf. Im. Proc.*, volume 3, pages 864–867, 2001.
- [3] I. Avcibas, I. Avcba, B. Sankur, and K. Sayood. Statistical evaluation of image quality measures. *Journal of Electronic Imaging*, 11:206–223, 2002.
- [4] W. Freeman, T. Jones, and E. Pasztor. Example-based super-resolution. *IEEE Computer Graphics and Applications*, 22(2):56–65, 2002.
- [5] A. Giachetti and N. Asuni. Fast artifact free image interpolation. In *Proc. BMVC 2008*, 2008.
- [6] K. Kamimura, N. Tsumura, T. Nakaguchi, Y. Miyake, and H. Motomura. Video super-resolution using texton substitution. In *ACM SIGGRAPH 2007 posters*, page 63, New York, NY, USA, 2007. ACM.
- [7] K. I. Kim and Y. Kwon. Example-based learning for single-image super-resolution. In *Proceedings of the 30th DAGM symp. on Patt. Rec.*, pages 456–465, Berlin, Heidelberg, 2008. Springer-Verlag.
- [8] X. Li and M. T. Orchard. New edge-directed interpolation. *IEEE Trans. on Image Proc.*, 10:1521–1527, 2001.
- [9] B. Morse and D. Schwartzwald. Image magnification using level-set reconstruction. In *Proc. IEEE Conf. Comp. Vis. Patt. Rec.*, volume 3, pages 333–340, 2001.
- [10] D. Su and P. Willis. Image interpolation by pixel level data-dependent triangulation. *Computer Graphics Forum*, 23, 2004.
- [11] J. Sun, Z. Xu, and H. Shum. Image super-resolution using gradient profile prior. In *CVPR08*, pages 1–8, 2008.
- [12] J. Sun, N. Zheng, H. Tao, and H. Shum. Image hallucination with primal sketch priors. In *Proceedings IEEE conf. on Comp. Vis. and Patt. Rec.*, volume 2, 2003.
- [13] Y. Tai, W. Tong, and C. Tang. Perceptually-inspired and edge-directed color image super-resolution. In *Proc. IEEE Conference on Comp. Vision and Patt. Recognition*, 2006.
- [14] N. Toronto, B. Morse, K. Seppi, and D. Ventura. Super-resolution via recapture and bayesian effect modeling. In *Proc. IEEE Conference on Comp. Vision and Pattern Recognition*, pages 2388–2395, 2009.
- [15] J. van Ouwkerk. Image super-resolution survey. *Image and Vision Computing*, 24:1039–1052, 2006.