Solving the inverse problem of image zooming using "self-examples"

Mehran Ebrahimi and Edward R. Vrscay

Department of Applied Mathematics
Faculty of Mathematics
University of Waterloo
Waterloo, Ontario, Canada N2L 3G1
m2ebrahi@uwaterloo.ca, ervrscay@uwaterloo.ca

Abstract. In this paper we present a novel single-frame image zooming technique based on so-called "self-examples". Our method combines the ideas of fractal-based image zooming, example-based zooming, and nonlocal-means image denoising in a consistent and improved framework. In Bayesian terms, this example-based zooming technique targets the MMSE estimate by learning the posterior directly from examples taken from the image itself at a different scale, similar to fractal-based techniques. The examples are weighted according to a scheme introduced by Buades *et al.* to perform nonlocal-means image denoising. Finally, various computational issues are addressed and some results of this image zooming method applied to natural images are presented.

1 Introduction

The idea of increasing the resolution of an image using its nonlocal self-similarities was initially proposed by the researchers in fractal image coding back in the 1990s – see, for example, [4, 17, 25, 27, 21]. In fractal image coding, one seeks to approximate a target image by the fixed point of a contractive, resolution-independent operator called the fractal transform. The essence of the fractal transform is to approximate smaller *range* subblocks of an image with modified copies of larger *domain* subblocks. Repeated action of the transform introduces finer details, implying that its fixed point can increase the resolution of images to any arbitrary degree. Because of the block coding involved in the construction of the fractal transform, however, the results are blocky and, in many ways, artificial.

A more recent series of novel *example-based methods* that take advantage of the self-similarity of images show great promise for the solution of various inverse problems [6, 7, 9, 14, 32]. In [6, 7], the authors demonstrated that their nonlocal-means (henceforth to be referred to as "NL-means") image denoising algorithm has the ability to outperform classical denoising methods. In this paper, we generalize their method to address the zooming problem. Similar to NL-means denoising, we compare neighborhoods in the image that will allow us to predict pixel values that reproduce local patterns and textures. Our algorithm compares the neighborhood information of each pixel to the neighborhoods of the same size

in the coarser scale over the entire image. Due to the reconstruction scheme, our results do not suffer from the blockiness generally inherent in fractal zooming. Moreover, the geometry of objects is well preserved for the test images presented.

In Section 2, we introduce the problem of image zooming, and present a brief introduction to the ideas of interpolation, regularization and classical inverse theory techniques, and example-based methods. In Section 3, we review the role of self-similarity in various inverse problems, namely fractal-based methods and example-based approaches. We explain how our example-based approach in this paper takes advantage of self-similarity across scales similar to fractal-based methods. In Section 4, we introduce our method, and discuss some computational issues in Section 5. Finally, some concluding remarks are presented in Section 6.

2 Some background on the inverse problem of image zooming

In the characterization of a visual image, the term "resolution" can be confusing since it involves a rather large number of competing terms and definitions. *Spatial resolution* refers to the spacing of pixels in an image. Many imaging devices, such as charged-coupled-device (CCD) cameras, consist of arrays of light detectors. A detector determines pixel intensity values depending upon the amount of light detected from its assigned area in the scene. The spatial resolution of images produced by these types of devices is proportional to the density of the detector array. In many applications, the imaging sensors have poor resolution output.

The process of producing a high-resolution image from a single lower-resolution and distorted (e.g., blurred, noisy) image is called *(single-frame) image zooming*.

2.1 The inverse problem of image zooming

We consider the following degradation model [8, 15, 16]

$$\mathbf{u} = \mathcal{H}\mathbf{f} + \mathbf{n}, \tag{1}$$

in which $\mathbf{u} \in l^2(\Omega)$ is the low-resolution $m \times n$-pixel observation, i.e.,

$$\Omega = [1, \ldots, m] \times [1, \ldots, n],$$

and $\mathbf{n} \in l^2(\Omega)$ is additive white, independent Gaussian noise with zero-mean and variance σ^2. $\mathbf{f} \in l^2(\Psi)$ is the high-resolution image to be recovered such that

$$\Psi = [1, \ldots, mz] \times [1, \ldots, nz],$$

where z is a positive integer. The operator $\mathcal{H} = \mathcal{S}\mathcal{B}$ is the composition of a blurring operator $\mathcal{B}$ followed by a downsampling operator $\mathcal{S}$ of factor z in each direction. More technically, in order to be consistent with the CCD-array model,

we assume that the blurring operator $\mathcal{B} : l^2(\Psi) \rightarrow l^2(\Psi)$ is the local averaging operator of length z, i.e., for any $(p, q) \in \Psi$,

$$(\mathcal{B}\mathbf{f})(p, q) = \frac{1}{z^2} \sum_{0 \leq p_1 < z, 0 \leq q_1 < z} \mathbf{f}(p + p_1, q + q_1). \tag{2}$$

Note that boundary conditions on $\mathbf{f}$ may be required so that $\mathcal{B}$ is well-defined. Also, the downsampling operator $\mathcal{S} : l^2(\Psi) \rightarrow l^2(\Omega)$ is defined for any image $\mathbf{I} \in l^2(\Psi)$ such that for any $(i, j) \in \Omega$,

$$(\mathcal{S}\mathbf{I})(i, j) = \mathbf{I}\Big((i-1)z + 1, (j-1)z + 1\Big). \tag{3}$$

The **problem of image zooming** can be viewed as the recovery of an estimate of $\mathbf{f}$, denoted by $\widetilde{\mathbf{f}}$, given the observation $\mathbf{u}$ and scaling factor z.

2.2 A brief introduction to various techniques

Image interpolation: The easiest and not necessarily most appealing approach to estimate $\mathbf{f}$ is to ignore the effect of noise $\mathbf{n}$ and the blurring operator $\mathcal{B}$, thereby assuming that $\mathbf{u}$ is provided by downsampling of $\mathbf{f}$, i.e., $\mathbf{u} = \mathcal{S}\mathbf{f}$. In this case, one could simply apply traditional interpolation algorithms on $\mathbf{u}$ to estimate $\mathbf{f}$, including *nearest neighbor or pixel replication*, *bilinear*, *bicubic*, *spline*, or *sinc interpolation*. In general, the resizing of an image does not translate into an increase in its resolution. Resizing should be accompanied by approximations to frequencies higher than those present in the observation $\mathbf{u}$ as well as at a higher signal-to-noise ratio. However some of these algorithms provide overly smoothed results around edges and fine details. The application of sharpening operators on these oversmoothed results does not usually recover the image details [8].

Regularization and classical inverse theory techniques: Single image zooming cannot recover high-frequency components that are lost during the low-resolution sampling process unless we include additional *a priori* information about the image. This can be normally done by using regularization. Regularization not only acts as an algebraic stabilizer in estimating the solutions of ill-posed inverse problems, but it may also improve the solutions of well-posed problems. Equivalently, in the Bayesian point of view, the probability density function (PDF) of the image, serving as the prior $p(\mathbf{f})$, plays the role of regularization. The most trivial form of regularization, known as Tikhonov regularization [30] that applies a uniform spatial smoothness on the outcome will over-smooth the edges. According to [15]: "Much of the progress made in the past two decades on inverse problems in image processing can be attributed to the advances in forming or choosing the way to practice the regularization. This includes, the regularization terms based on variational methods [29], edge-preserving regularization terms formed using various partial differential equations (PDE) methods [33], as well as sparsity of transform coefficients (e.g. wavelet) [11]." *The effectiveness of single-image zooming techniques is always limited to the prior used in the reconstruction process.*

Example-based methods: A very attractive approach in solving imaging inverse problems is to exploit examples in defining the PDF of the image instead of intuitively defining a regularization term. There are various ways to apply examples in inverse problems as comprehensively described in [15]: Examples can be used either (a) directly in the reconstruction procedure [9, 14, 19, 20, 26, 32], (b) in tuning the parameters of previously defined regularization expressions [23, 28, 35] or (c) in an alternating combination of these two methods [3, 10, 15, 19, 20]. The method described in this paper corresponds to making use of examples *directly* in the reconstruction procedure, i.e., method (a).

3 The role of self-similarity in various inverse problems

3.1 Fractal-based methods

The original motivation of fractal coding was image compression [4, 25]. Most methods rely upon the method of fractal block coding, where the subimage $\mathbf{u}|_{R_i}$ supported on a smaller *range* block R_i is approximated by a spatially contracted and greyscale-modified copy of the subimage $\mathbf{u}|D_j$ supported on a larger *domain* block D_j. (Of course, one has to search for such a domain block.) In this way, fractal coding exploits the local (affine) self-similarity of image patches at two different scales. Some variations and improvements of fractal coding include: variable partitioning schemes [17], redundant representation and multiple-domain methods [5, 34, 2], lapped approximations [24], and fractal coding on wavelet coefficient trees (IFSW) [18, 31]. More recently, the ability of fractal-coding scheme to solve other inverse problems in imaging has been investigated. For example, it has been observed that fractal-based methods have denoising capabilities [22, 1]. Also, due to the resolution independence nature of the fractal-transform operator, interpolation algorithms called "fractal zoom" have been developed in the literature [17, 21, 25, 27].

Some recent investigations [1, 2] have shown that images generally possess a great deal of local (affine) self-similarity: Given a subimage $\mathbf{u}|_{R_i}$ there are often a good number of domain blocks D_j whose subimages $\mathbf{u}|_{D_j}$ approximate it as well as the "best" domain block. This feature, which never seems to have been quantified previously, accounts for the rather small degradations that are experienced when the size of the *domain pools* – the domain blocks D_j to be examined – is decreased.

As mentioned earlier, however, the images produced by most fractal-based methods, including "fractal zoom," suffer from blockiness. Another major drawback of traditional fractal-based methods is that their output is quite restricted, i.e., no prior knowledge, extra regularization method or tuning (regularization) parameters can be combined with these methods. Recently, we have examined various possibilities to address this problem in [12, 13]. In this paper, we shall combine some of the ideas in the above works in a natural and complementary fashion.

3.2 Self-Similarity in various example-based approaches

In the work of [9, 14, 32] on texture synthesis, and inpainting (filling in holes), examples are taken from the image *itself*.

Another important example-based approach, nonlocal-means (NL-means) image denoising [6, 7] addresses the denoising problem using examples from the noisy image *itself* at the *same scale* with a Gaussian-type weighting scheme. The authors demonstrated that their algorithm has the ability to outperform classical denoising methods, including Gaussian smoothing, Wiener filter, TV filter, wavelet thresholding, and anisotropic diffusion. Since we shall generalize this algorithm to address the zooming problem, it is necessary to review the algorithm below.

NL-means image denoising [6, 7]: Consider the following image denoising problem

$$\mathbf{u} = \mathbf{f} + \mathbf{n}.$$

For any $x \in \Omega$ define the denoising formula

$$NL\big(\mathbf{u}(x)\big) = \widetilde{\mathbf{f}}(x) = \frac{1}{C(x)} \sum_{y \in \Omega} w(x,y)\mathbf{u}(y), \text{ such that} \tag{4}$$

$$w(x,y) = \exp\Big(- \frac{\| \mathbf{u}(\mathcal{N}^d\{x\}) - \mathbf{u}(\mathcal{N}^d\{y\}) \|^2_{2,a}}{h^2}\Big), \text{ and} \tag{5}$$

$$C(x) = \sum_{y \in \Omega} w(x,y), \tag{6}$$

where the expressions $\mathcal{N}^d\{\ldots\}$ and $\| \, . \, \|^2_{2,a}$ are defined in the following way.

Neighborhoods: For any point in the domain of observation $(i,j) \in \Omega$, define

$$\mathcal{N}^d\big\{(i,j)\big\} = \Big\{(i+i_1, j+j_1) \mid (i_1,j_1) \in \mathbb{Z}^2, \max\{|i_1|,|j_1|\} \leq d\Big\}. \tag{7}$$

Gaussian-semi-norm: For an image $\mathbf{I}$ the Gaussian-weighted-semi-norm $\| \, . \, \|_{2,a}$ is defined in terms of the l^2 norm as

$$\| \mathbf{I} \|_{2,a} = \| \, G_a \star I \, \|_2$$

in which G_a is a two-dimensional Gaussian kernel of standard deviation a.

The idea of the NL-means algorithm is that given a discrete noisy image $\mathbf{u}$, the estimated denoised value $NL\big(\mathbf{u}(x)\big)$ is computed as a weighted average of all the pixels in the image, $\mathbf{u}(y)$, where the weights $w(x,y)$ depend on the similarity of neighborhoods of the pixels x and y, and w is a decreasing function of the weighted Euclidean distance of the neighborhoods. The parameter h acts as a degree of filtering. It controls decay of the exponential function and therefore the decay of the weights as a function of the Euclidean distances.The NL-means not only compares the grey level in a single point but the geometrical configuration in a whole neighborhood [6, 7].

3.3 From same-scale towards across-scale self-examples

The NL-means algorithm mentioned above [6, 7] along with the work reported in [14] use self-examples, i.e., they take examples from the image itself. However, in all cases, *the examples are taken at the same scale*. This clearly represents a major difference between these example-based methods and fractal-based methods.

In the zooming algorithms proposed below, we take advantage of the richness of the NL-means algorithm (and thereby overcome the issue of blockiness) yet remain consistent with fractal-based methods by *comparing patches across scales*.

4 Image zooming algorithm using self-examples

In this section we introduce our notations, and then formulate a natural extension of the NL-means denoising method, to be denoted as Algorithm 1 below. Our main algorithm (image zooming algorithm using self-examples) will be a special case of Algorithm 1, in which the example image is precisely the input image. This formulation will detect across-scale similarities.

Example image: The notion of *example image* is denoted by $\mathbf{v} \in l^2(\Phi)$, where Φ is the $k \times l$ pixel lattice defined by

$$\Phi = [1, \ldots, k] \times [1, \ldots, l].$$

Lattice covers: Recall that

$$\Psi = [1, \ldots, mz] \times [1, \ldots, nz]$$

is the lattice on which the high-resolution image $\mathbf{f}$ was defined. Let $\mathcal{P}$ be the partition of the lattice Ψ composed of non-overlapping square blocks of size $z \times z$. Also, define $\mathcal{Q}$ to be the set of all $z \times z$ square blocks on the lattice Φ. We allow the elements of $\mathcal{Q}$ to overlap, implying that $\mathcal{Q}$ does not necessarily form a partition of Φ.

Lattice mapping: Let the mapping $\mathcal{D} : \mathcal{P} \to \Omega$ be defined in the following way. For any $z \times z$ block $x \in \mathcal{P}$ where

$$x = \Big\{(x_1, x_2) \in \mathbb{N}^2 \mid \lceil \frac{x_1}{z} \rceil = i, \lceil \frac{x_2}{z} \rceil = j, \text{ for some fixed } (i, j) \in \Omega\Big\},$$

define $\mathcal{D}_x = (i, j)$.

Extended z-neighborhoods: For any block y in example-image-cover $y \in \mathcal{Q}$ define

$$\mathcal{N}_z^d\big\{y\big\} = \Big\{(y_1+i_1, y_2+i_2) \mid (y_1, y_2) \in y, (i_1, i_2) \in \mathbb{Z}^2, \max\{|i_1|, |i_2|\} \leq dz\Big\}. \quad (8)$$

The following is a natural extension of the NL-means algorithm. In Bayesian terms the algorithm is an MMSE estimation by learning the posterior directly from the examples taken from an *example-image* $\mathbf{v}$.

Algorithm 1. An example-based zooming algorithm The approximation of $\mathbf{f}$, denoted by $\widetilde{\mathbf{f}}$, given the observation $\mathbf{u}$, using *example image* $\mathbf{v}$, is computed in the following way. For any $x \in \mathcal{P}$,

$$\widetilde{\mathbf{f}}_{(\mathbf{u},\mathbf{v},z,h,d,a)}(x) = \frac{1}{C(x)} \sum_{y \in \mathcal{Q}} w(x,y)\mathbf{v}(y) \tag{9}$$

where,

$$w(x,y) = \exp\Big(- \frac{\| \mathbf{u}(\mathcal{N}^d\{\mathcal{D}x\}) - \mathcal{H}\big(\mathbf{v}(\mathcal{N}_z^d\{y\})\big) \|_{2,a}^2}{h^2} \Big) \tag{10}$$

and,

$$C(x) = \sum_{y \in \mathcal{Q}} w(x,y). \tag{11}$$

Proposition 1. *NL-means image denoising is equivalent to Algorithm 1 for the case that* $z = 1$ *and* $\mathbf{v} = \mathbf{u}$*, i.e.,*

$$NL\big(\mathbf{u}(x)\big) = \widetilde{\mathbf{f}}_{(\mathbf{u},\mathbf{u},1,h,d,a)}(x). \tag{12}$$

Proof. If $z = 1$ and $\mathbf{v} = \mathbf{u}$ then $\Omega = \Psi = \Phi = \mathcal{P} = \mathcal{Q}$, $\mathcal{D}x = x$, and $\mathcal{H}$ becomes the identity operator. Furthermore, in this case $\mathcal{N}_1^d\big\{y\big\} = \mathcal{N}^d\big\{y\big\}$. Substituting the corresponding expressions in Algorithm 1 leads to the equivalent expression of the NL-means algorithm.

The benefit of our notation and the introduction of Algorithm 1 above is that a simple substitution, namely, $\mathbf{v} = \mathbf{u}$, will lead to our proposed technique of "Image zooming using self-examples," the main algorithm of this paper.

Algorithm 2. Image zooming using self-examples For a given observation $\mathbf{u}$ and a zooming factor z apply Algorithm 1, using the observation in place of the example image as well, i.e., $\mathbf{v} = \mathbf{u}$, and estimate $\widetilde{\mathbf{f}}_{(\mathbf{u},\mathbf{u},z,h,d,a)}$.

Again it can be observed that Algorithm 2 is a special case of Algorithm 1, naturally by the constraint $\mathbf{v} = \mathbf{u}$, yet still an extension of the NL-means denoising if $z = 1$. Similar to NL-means denoising, the main algorithm, from a Bayesian viewpoint, targets the MMSE estimate by learning the posterior directly from the examples taken from the image itself at a different scale. As mentioned above, this is similar to fractal-based approaches.

5 Numerical considerations and results

5.1 Various numerical considerations

Several aspects should be carefully considered in the implementation of our algorithm. It may be necessary to apply appropriate boundary conditions, for example, when encountering neighborhoods that fall out of the image domain or

applying the operator $\mathcal{H}$ on patches. In all of our experiments we have applied symmetric boundary conditions. Furthermore, the adjustment of the parameter h varies the degree of filtering in a manner quite similar to that of NL-means denoising. A large h-value will blur the output result; a very small h-value will amplify any image noise in the output.

It should also be noticed that if finite precision is employed in our computations, then if h is sufficiently small, then all weighting terms in Equation (10) will be computed as zero. In other words, it is possible that for some $x \in \mathcal{P}$ and for some small pre-chosen parameter $h > 0$ the weights are computed as $w(x, y) = 0$ for all values of $y \in \mathcal{Q}$. Computationally, this is not an issue in the NL-means image denoising because there is at least one patch – the neighborhood of x itself – that makes the exponent equal to 1 regardless of the value of $h > 0$. We can bypass this problem by defining $w(x, y) = 1_S(y)$, at the problematic points $x \in \mathcal{P}$, where $1_S(y)$ denotes the indicator function of y over the set S defined as follows,

$$S = \left\{ y_s \in \mathcal{Q} \mid y_s = \arg\min_y \left(\| \mathbf{u}(\mathcal{N}^d\{\mathcal{D}x\}) - \mathcal{H}\big(\mathbf{v}(\mathcal{N}_z^d\{y\})\big) \|_{2,a} \right) \right\}. \quad (13)$$

This is equivalent to averaging over all $\mathbf{v}(y)$ for which y is a minimizer of the similarity distance.

5.2 Computational results

In this section we present some results of our computational experiments. In all experiments we have chosen a zoom factor $z = 2$, and 3×3 neighborhoods as well as 6×6 extended-neighborhoods, i.e., $d = 3$. We also have chosen $G_a = \delta$, i.e., simply ignoring the effect of Gaussian smoothing. As well, we have varied the parameter h with respect to each image individually.

Figure 1 indicates the result of an experiment of Algorithm 1 to show that using irrelevant examples (middle column) may lead to irrelevant outputs (right column). In Figure 2, we present the outputs of Algorithm 2 for a noiseless (top left) and noisy (bottom left) image. It can be seen that the algorithm performs denoising in parallel with zooming. For purposes of comparison, the results of pixel replication and bilinear interpolation are also presented.

In rows 1 and 3 of Figure 3, we present the result of Algorithm 2 when applied to a simple circular shape in both noiseless and noisy situations, respectively. Below each of these rows are displayed the Fourier spectra of each image. In the bottom two rows of Figure 3, we show the effect of Algorithm 2 on a Gaussian white noise sample, along with corresponding Fourier spectra. The value of parameter h was set in a way that the variance of noise was three-quarter of the variance of the input noise. However, the output is still noise-like, as opposed to the results of bilinear and pixel replication. It can be seen that the Fourier spectrum of the output is also distributed rather uniformly over the frequency domain, as opposed to the pixel replication and bilinear interpolation cases.

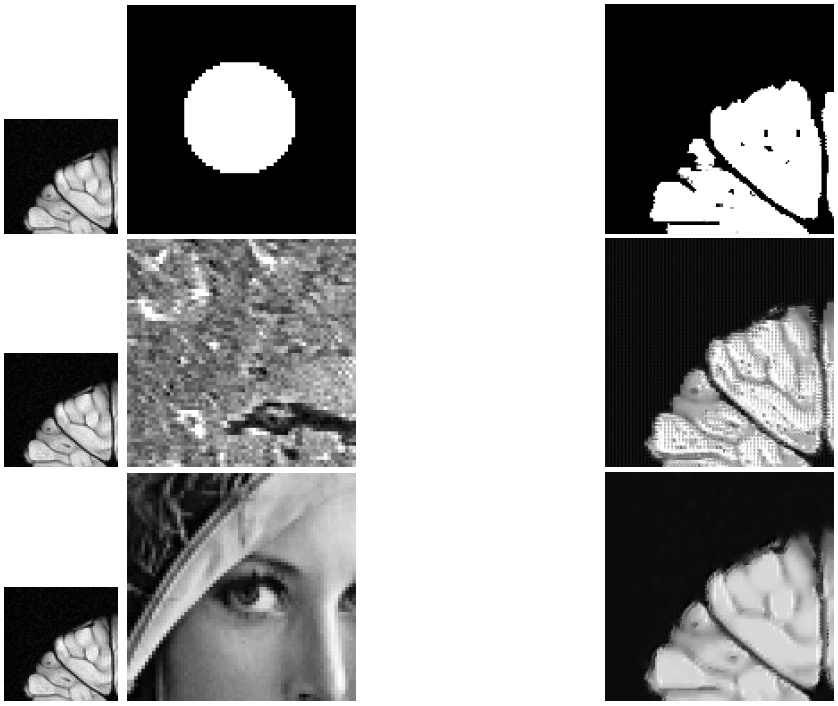

Fig. 1. Original, Example and the image reconstructed using Algorithm 1, $h = 0.1$ in all three cases

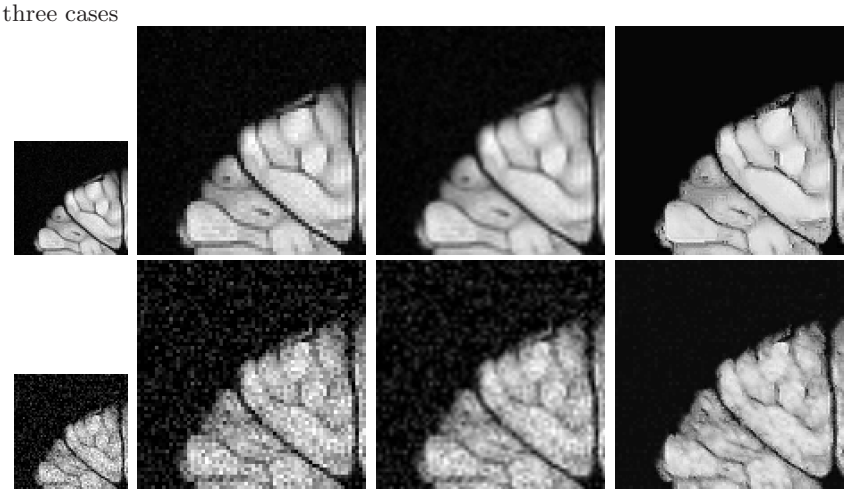

Fig. 2. Original , Pixel replication, Bilinear, Self-examples. In the first row $h = 0.05$ is applied while in the second row the standard deviation of noise $\sigma = 0.1$ and $h = 0.15$.

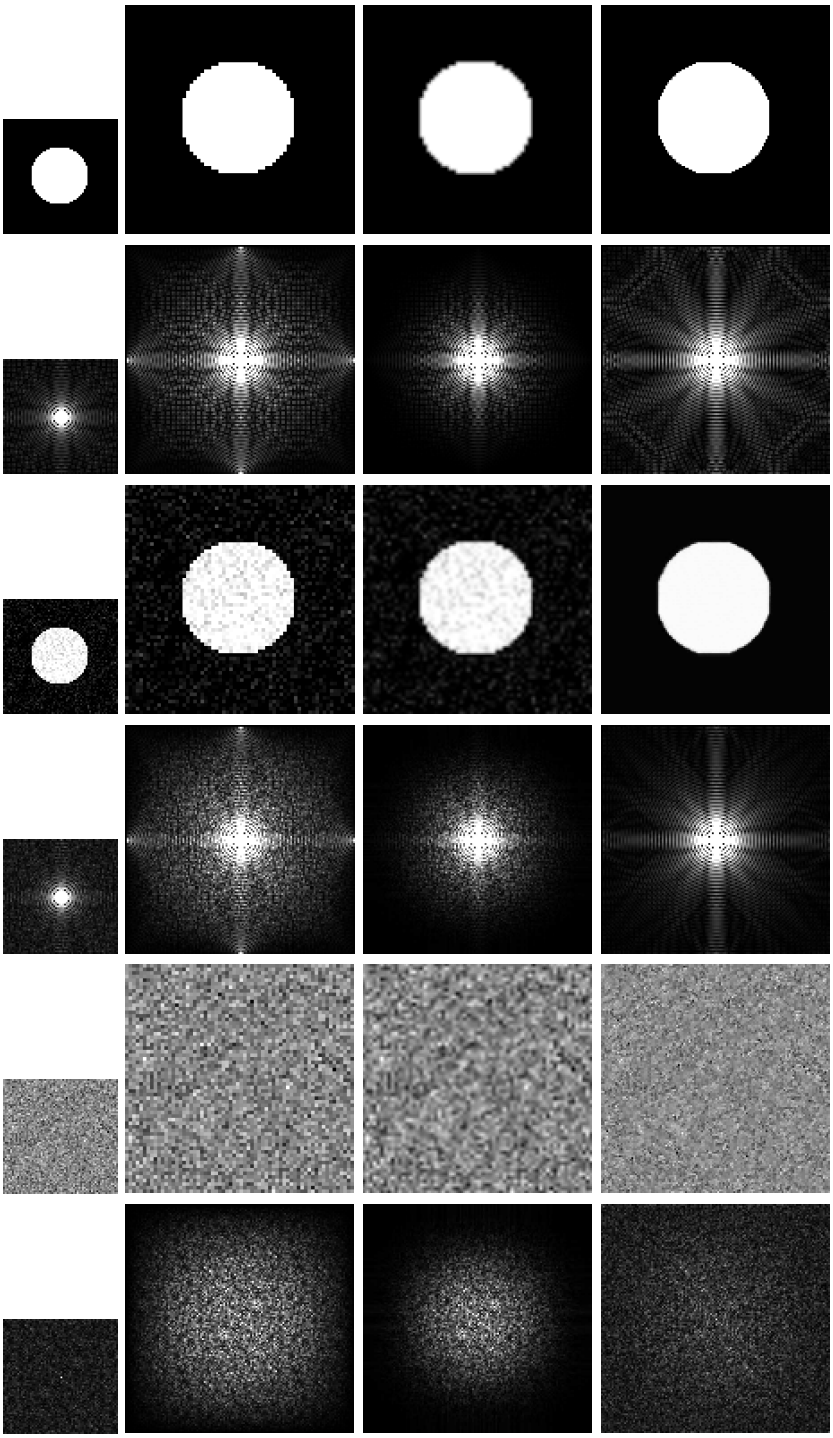

Fig. 3. A comparison of the image zooming using self-examples, Algorithm 2, with other methods of zooming. Three input images are considered: Circular region (row 1), noisy circular region (row 3) with the noise of standard deviation $\sigma = 0.1$, pure noise image (row 5). The Fourier spectra of all images are also shown (rows 2,4 and 6). Starting at left: input image, zooming with pixel replication, zooming with bilinear interpolation, zooming with Algorithm 2. In rows 1,3,5 the value of h in the experiment are respectively 0.01, 0.5, and 0.05.

6 Conclusions

In this paper we have presented a novel single-frame image zooming method of self-examples, explaining how it combines the ideas of fractal-based image zooming, example-based zooming and nonlocal-means image denoising. Our framework implicitly defines a regularization scheme which exploits the examples taken from the image itself at a different scale in order to achieve image zooming. The method essentially extends the NL-means image denoising technique to image zooming problem. Various computational issues and results were also presented, showing that frequency domain extrapolation is in fact possible with this method. The issue of parameter selection, e.g., the filter parameter h, is an important issue which has not been addressed in sufficient detail this paper. We are currently investigating adaptive parameter schemes that may be able to achieve superior results.

Acknowledgments. This research has been supported in part by the Natural Sciences and Engineering Research Council of Canada, in the form of a Discovery Grant (ERV). M.E. is supported by an Ontario Graduate Scholarship.

References

1. S.K. Alexander, *Multiscale methods in image modelling and image processing*, Ph.D. Thesis, Dept. of Applied Mathematics, University of Waterloo, 2005.
2. S.K. Alexander, E.R. Vrscay, and S. Tsurumi. *An examination of the statistical properties of domain-range block matching in fractal image coding*, Preprint, 2007.
3. S. Baker, and T. Kanade, *Limits on super-resolution and how to break them*, IEEE Trans. Patt. Analysis and Mach. Intel. Vol. 24, No. 9, pp. 1167-1183, 2002.
4. M.F. Barnsley, *Fractals Everywhere*, Academic Press, New York, 1988.
5. D.J. Bone, *Orthonormal fractal image encoding using overlapping blocks*, Fractals, vol. 5 (Supplementary Issue), pp. 187-199, Apr. 1997.
6. A. Buades, B. Coll, and J.M. Morel, *A nonlocal algorithm for image denoising*, IEEE International conference on Computer Vision and Pattern Recognition (CVPR), Vol. 2, pp. 6065, San-Diego, California, June 20-25, 2005.
7. A. Buades, B. Coll, and J.M. Morel, *A review of image denoising algorithms, with a new one*, SIAM Journal on Multiscale Modeling and Simulation (MMS), Vol. 4, No. 2, pp. 490-530, 2005.
8. S. Chaudhuri, *Super-resolution imaging*, Boston, MA, Kluwer, 2001.
9. A. Criminisi, P. Perez., and K. Toyama, *Region filling and object removal by exemplar-based image inpainting*, IEEE Trans. on Image Proc., Vol. 13, No. 9, pp. 1200-1212, 2004.
10. D. Datsenko, and M. Elad, *Example-based single document image superresolution: A global MAP approach with outlier rejection*, to appear in the Journal of Mathematical Signal Processing, 2006.
11. D.L Donoho, and I.M. Johnstone, *Ideal spatial adaptation by wavelet shrinkage*, Biometrika Vol. 81 No. 3, pp. 425-455, 1994.
12. M. Ebrahimi and E.R. Vrscay, *Fractal image coding as projections onto convex sets*, Lecture Notes in Computer Science, Volume 4141, Book: Image analysis and Recognition, pp. 493-506, Springer Berlin/Heidelberg, 2006.

13. M. Ebrahimi and E.R. Vrscay, *Regularized fractal image decoding*, Proceedings of CCECE '06, Ottawa, Canada, pp.1933-1938, May 7-10, 2006.
14. A.A. Efros, and T.K. Leung, *Texture synthesis by non-parametric sampling*, IEEE International Conference on Computer Vision (ICCV), Corfu, Greece, pp. 1033-1038, September 20-25, 1999.
15. M. Elad and D. Datsenko, *Example-Based Regularization Deployed to Super-Resolution Reconstruction of a Single Image*, to appear in The Computer Journal.
16. M. Elad, and A. Feuer, *Restoration of a single superresolution image from several blurred, noisy, and undersampled measured images*, IEEE Trans. on Image Proc., vol. 6, no. 12, pp. 1646-1658, 1997.
17. Y. Fisher, Ed., *Fractal image compression, theory and application*, New York, Springer-Verlag, 1995.
18. B. Forte, and E.R. Vrscay, *Theory of generalized fractal transforms*, in Fractal image encoding and analysis., Y. Fisher, Ed. New York: Springer-Verlag, 1998.
19. W.T. Freeman, E.C. Pasztor, and O.T. Carmichael, *Learning low-level vision*, Int. Journal Of Computer Vision, Vol. 40, No. 1, pp. 25-47, 2000.
20. W.T. Freeman, T.R. Jones, and E.C. Pasztor, *Example-based super-resolution*, IEEE Comp. Graphics And Appl., Vol. 22, No. 2, pp. 56-65, 2002.
21. M. Gharavi-Al., R. DeNardo, Y. Tenda, and T.S. Huang, *Resolution enhancement of images using fractal coding*, Visual Communications and Image Processing, vol. 3024 of SPIE Proceedings, (San Jose, CA, USA), pp. 1089-1100, Feb. 1997.
22. M. Ghazel, G. Freeman, and E.R. Vrscay, *Fractal image denoising*, IEEE Trans. on Image Proc. **12**, no. 12, 1560-1578, 2003.
23. E. Haber, E. and L. Tenorio, *Learning regularization functionals*, Inverse Problems, Vol. 19 pp. 611-626. 2003.
24. H. Ho, W. Cham, *Attractor Image Coding using Lapped Partitioned Iterated Function Systems*, icassp, p. 2917, IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP'97), Volume 4, 1997.
25. N. Lu, *Fractal Imaging*, Academic Press, NY, 1997.
26. R. Nakagaki, and A.K. Katsaggelos, *VQ-based blind image restoration algorithm*, IEEE Trans. On Image Proc., Vol. 12, No. 9, pp. 1044-1053, 2003.
27. E. Polidori and J.-L. Dugelay, *Zooming using iterated function systems*, Fractals, vol. 5 (Supplementary Issue), pp. 111123, Apr. 1997.
28. S. Roth, S. and M.J. Black, *Fields of experts: A framework for learning image priors*, IEEE Conference on Computer Vision and Pattern Recog. (CVPR), Vol. 2, pp. 860-867, San-Diego, California, June 20-25, 2005.
29. L. Rudin, S. Osher, and E. Fatemi, *Nonlinear total variation based noise removal algorithms*, Physica D, Vol. 60, pp. 259-268, 1992.
30. A.N. Tikhonov, and V.A. Arsenin, *Solution of Ill-posed Problems*, Winston & Sons, Washington, 1977.
31. E.R. Vrscay, *A generalized class of fractal-wavelet transforms for image representation and compression*, Can. J. Elect. Comp. Eng. vol. 23, no. 1-2, pp. 69-84, 1998.
32. L.Y. Wei, and M. Levoy, *Fast texture synthesis using tree-structured vector quantization*, Proc. of SIGGRAPH, pp. 479-488, New Oleans, Louisiana, 2000.
33. J. Weickert, *Anisotropic Diffusion in Image Processing*, ECMI Series, Teubner, Stuttgart, 1998.
34. W. Xu, D. Fussell, *IFS coding with multiple DC terms and domain blocks*, Citeseer article 185324, available at: http://citeseer.ist.psu.edu/185324.html
35. S.C. Zhu and D. Mumford, *Prior learning and Gibbs reaction-diffusion, IEEE Trans. on Patt. Analysis and Machine* Intel., Vol. 19, No. 11, pp. 1236-1250, 1997.