

IMAGE UPSAMPLING VIA SPATIALLY ADAPTIVE BLOCK-MATCHING FILTERING

Aram Danielyan, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian

Department of Signal Processing, Tampere University of Technology
P.O. Box 553, 33101, Tampere, Finland
web: www.cs.tut.fi/~comsens email: firstname.lastname@tut.fi

ABSTRACT

In this paper we present a novel wavelet-domain image upsampling algorithm based on iterative spatially adaptive filtering. A high-resolution image is reconstructed by alternating two procedures: spatially adaptive filtering and projection on the observation-constrained subspace. The Block-Matching and 3D filtering (BM3D) [3] technique is used to suppress ringing, and reconstruct missing wavelet detail coefficients. The BM3D algorithm exploits the local image statistics collected from similar blocks to extract local and non-local image features by 3D transform-domain shrinkage. It results in high-quality upsampled images, with sharp edges and practically no artifacts.

1. INTRODUCTION

Image upsampling is an intrinsically ill-posed problem. To solve it, one has to apply some constraints. The choice of the constraints and of the interpolation methods depends on assumptions about how the low-resolution image was obtained. Different approaches were suggested for different observation models. Let us mention few among the more recently published methods.

Sparsity. The method presented in [15] by Mueller et al., utilizes sparsity in a contourlet transform domain as a regularizing condition. The algorithm iteratively alternates projecting to the wavelet approximation subspace defined by the low-resolution image with thresholding the coefficients in contourlet domain. While performing well along edges, this method introduces significant ringing artifacts due to the length of the filters used in the contourlet transform.

Cross-scale correlation. Liu et al. [13] proposed a non-iterative algorithm based on high-pass subband estimation from the interpolated Laplacian image pyramid (WLPSR - wavelet Laplacian pyramid super resolution). To compensate the ringing, cycle-spinning is performed. Several low-resolution images are generated by spatial shifting and down-sampling from high-resolution image estimate obtained by WLPSR. These low-resolution images are reconstructed again with the WLPSR and fused into the final estimate by averaging.

Learning. The alternative approach based on learning by examples was introduced by Freeman et al. [7], where unknown high-frequency data is “learned” from a training set of high-resolution training images. A similar algorithm was proposed by Jiji et al. [8]. It is based on learning in the wavelet domain, where unknown wavelet coefficients at the finer scales are “learned” from the fine scale coefficients of the training images. A modification of this algorithm using contourlet learning [9] was also presented. To upsample textures, Li [11] proposed to use texture descriptive statistics in wavelet domain as parametric constraints, where parameters are estimated by learning from high-resolution images. The algorithm uses patch-based learning, combining the patches via Bayesian fusion.

This work was supported by the Academy of Finland (application no. 213462, Finnish Programme for Centres of Excellence in Research 2006-2011, and application no. 118312, Finland Distinguished Professor Programme 2007-2010).

Non-local approximations. The methods in this group exploit the idea of non-local averaging introduced by Buades et al. [1]. The zooming technique by Ebrahimi and Vrscay [5] combines learning from examples taken across different scales with the weighting scheme used in the denoising algorithm from [1]. Luong et al. [14] proposed a method where the found similar blocks are interpolated, registered with subpixel precision and finally fused in a median estimate. The additional TV regularization step is performed in order to denoise or deblur the upsampling data.

In this work, we propose a new upsampling algorithm based on iterative spatially adaptive filtering. This filtering has been successfully applied to many image reconstruction problems such as denoising [3] and compressed sensing [6]. There it has been shown that this filtering, being of nonparametric nature, produces results which are highly competitive with respect to the state-of-the-art algorithms developed on the base of a global optimization.

For the upsampling task considered in this paper, we use the Block Matching and 3D filtering (BM3D) algorithm [3]. The BM3D algorithm exploits the local image statistics collected from similar blocks. Yet, the filter is non-local, as the collected blocks can be at different spatial locations. The local and non-local image features are extracted simultaneously by the so-called *collaborative filtering*, which is realized by 3D transform-domain shrinkage. Thus, the filter is highly sensitive to the image details, and stable versus noisy components of the data.

In what follows, we limit ourself to the wavelet-domain observation model utilized by [15], for which the upsampling problem can be formalized as follows.

Let us assume that the *low-resolution image* y_{low} of size $n_1 \times n_2$ is obtained from the (coarsest) approximation subband $\Omega_{L^m L^m} = \Omega$ of an m -stage orthonormal wavelet decomposition \mathcal{W}_m of a *higher-resolution* original image y of size $2^m n_1 \times 2^m n_2$ as

$$y_{\text{low}} = 2^{-m} \mathcal{W}_m(y)|_{\Omega},$$

where $|_{\Omega}$ denotes the restriction to the subdomain Ω (also of size $n_1 \times n_2$) and the scaling factor 2^{-m} ensures that the means of y and y_{low} are the same. The problem is to reconstruct y from y_{low} . Clearly, any good estimate \hat{y} of y must have its approximation subband equal $2^m y_{\text{low}}$. Under this restriction, the estimates constitute an affine subspace $W_{y_{\text{low}}}$ of codimension $n_1 n_2$ in a $2^{2m} n_1 n_2$ -dimensional linear space W : $W_{y_{\text{low}}} = \{y \in W : \mathcal{W}_m(y)|_{\Omega} = 2^m y_{\text{low}}\}$. The obvious minimum ℓ^2 -norm estimate $\hat{y}_{\ell^2} = \arg \min_{y \in W_{y_{\text{low}}}} \|y\|_2$ is obtained when all the wavelet detail

coefficients (defined on the complementary Ω^c of Ω) are zero, i.e. $\hat{y}_{\ell^2} = \mathcal{W}_m^{-1}(2^m \mathcal{U}_m\{y_{\text{low}}\})$, where \mathcal{U}_m is a zero-padding operator, $\mathcal{U}_m(A) = \begin{bmatrix} A & 0 \\ 0 & 0 \end{bmatrix}$, that expands an input matrix A of size $\tilde{n}_1 \times$

\tilde{n}_2 to the size $2^m \tilde{n}_1 \times 2^m \tilde{n}_2$. Note that $\hat{y}_{\ell^2} = \mathcal{W}_m^{-1}(\mathcal{W}_m(y) \chi_{\Omega})$, where χ_{Ω} is the characteristic (indicator) function of Ω , $\chi_{\Omega}(i, j) = 1 \forall (i, j) \in \Omega$ and $\chi_{\Omega}(i, j) = 0 \forall (i, j) \notin \Omega$. However, the estimate obtained in this way would suffer from heavy ringing artifacts and blurred edges.




Figure 1: Iteration scheme

Our iterative upsampling process is initialized from the minimum ℓ^2 -norm solution \hat{y}_{ℓ^2} . At each iteration, the adaptive filtering is applied to the current estimate and then the filtered estimate is projected on the subspace defined by the low-resolution image y_{low} . The general scheme of our algorithm is essentially based on our previously proposed iterative procedure for compressed sensing image reconstruction¹ [6]. While, in the iterative alternation of the constraints, it is similar to the interpolation algorithm from [15], the block-matching and collaborative filtering procedure can be viewed as a self-learning, where each block “learns” or “tries to extract” missing information from neighboring similar blocks, thus resembling [8] or [5].

We demonstrate the effectiveness of the proposed approach in a number of experiments. The obtained results show a significant improvement in PSNR and SSIM [17] over some of the best methods known in the field.

2. THE ALGORITHM

2.1 Iterative system

One-stage upsampling. Given y_{low} , we define a sequence of estimates $\hat{y}^{(k)}$, $k = 0, 1, 2, \dots$, using the following iterative system:

$$\begin{cases} \hat{y}^{(0)} = \hat{y}_{\ell^2} = \mathcal{W}_m^{-1}(2^m \mathcal{U}_m \{y_{\text{low}}\}), \\ \hat{y}^{(k)} = \mathcal{W}_m^{-1}(2^m \mathcal{U}_m \{y_{\text{low}}\} + \mathcal{W}_m(\Phi(\hat{y}^{(k-1)}), \sigma_k))(1 - \chi_{\Omega}), \end{cases} \quad (1)$$

where Φ is a spatially adaptive filter, σ_k is a parameter controlling the strength of this filter, and $1 - \chi_{\Omega} = \chi_{\Omega^c}$ is the characteristic function of the complementary of Ω . In other words, at each iteration in (1), we filter the image $\hat{y}^{(k-1)}$ obtained from the previous iteration, perform wavelet decomposition, substitute the $n_1 \times n_2$ approximation coefficients with those defined by y_{low} (considered as true information known about original image y), and take an inverse wavelet transform to obtain $\hat{y}^{(k)}$. The flowchart of the system (1) is presented in Figure 1. The iteration process stops when the distance between $\hat{y}^{(k)}$ and $\hat{y}^{(k-1)}$ in some metric becomes less than certain threshold δ_0 , or if the maximum number of iterations k_{final} is reached.

We remind that the value of m determines the total number of decomposition levels in the wavelet transform \mathcal{W}_m . Because the size of the approximation subband Ω (obtained after the m -th decomposition) is fixed to $n_1 \times n_2$ (equal to the size of y_{low}), the value of m determines also the size of the high-resolution images. Therefore, the above scheme can be used to upsample y_{low} to any size $2^m n_1 \times 2^m n_2$, for an arbitrary $m \in \mathbb{N}$.

m -stage progressive upsampling. For upsampling by a factor 2^m when $m > 1$, an alternative scheme can be suggested. Instead of upsampling the image 2^m times at once, we upsample it progressively m times, each time (we call it *stage*) by the factor of 2. In this case the output from previous stage \hat{y}_{j-1} becomes the input $\hat{y}_j^{(0)}$

¹Note however that here our notation differs from that in [6], where y and \hat{y} were used to denote $2D$ transforms of the image and its estimate, respectively.

for the next one, where stage numbers are denoted by the subscript $j = 1, \dots, m$. In this way, the recursion equation takes the following form:

$$\begin{cases} \hat{y}_0 = y_{\text{low}}, \\ \hat{y}_j^{(0)} = \mathcal{W}_1^{-1}(2\mathcal{U}_1 \{\hat{y}_{j-1}\}), \\ \hat{y}_j^{(k)} = \mathcal{W}_j^{-1}(2^j \mathcal{U}_j \{y_{\text{low}}\} + \mathcal{W}_j(\Phi(\hat{y}_j^{(k-1)}), \sigma_k))(1 - \chi_{\Omega}), \end{cases} \quad (2)$$

where the superscript k corresponds to the iteration count inside each stage. We remark that this m -stage upsampling is *not* a recursion of m “one-stage upsamplings with factor 2”, because in (2) the projection is always made onto the initial $\mathcal{W}_{y_{\text{low}}}$ defined by y_{low} (and not by $\hat{y}_j^{(0)}$).

2.2 The filter

We use the BM3D denoising algorithm as the spatially adaptive filter Φ . Its detailed description can be found in [3]. In brief, the filter works as follows.

1. *Block-wise estimates.* Processing the image in sliding-window manner, for each block:
 - (a) *Grouping.* Find blocks that are similar to the currently processed one and then stack them together in a 3D array (group).
 - (b) *Collaborative hard-thresholding.* Apply a 3D transform to the formed group, attenuate the noise by hard-thresholding of the transform coefficients, invert the 3D transform to produce estimates of all grouped blocks, and return the estimates of the blocks to their original positions.
2. *Aggregation.* Compute the estimate of the output image by weighted averaging all of the obtained block-wise estimates that are overlapping.

Due to the similarity between the grouped blocks, the transform can achieve a highly sparse representation of the true signal so that the noise or small distortions can be well separated by shrinkage. In this way, the collaborative filtering reveals even the finest details shared by grouped fragments and at the same time it preserves the essential unique features of each individual fragment.

For the purposes of this work, we do not perform the final collaborative Wiener filtering stage of the original BM3D denoising algorithm [3].

Here, the parameter σ_k is used in place of the standard-deviation of the noise. This parameter controls the similarity tolerance of the grouping and the collaborative hard-thresholding strength. In order to prevent smearing of the small details the sequence $\{\sigma_k\}_{k=0,1,\dots}$ should be decreasing with the progress of the iterations.

3. EXPERIMENTAL RESULTS

We performed two sets of experiments: first, we assess both objective and subjective quality with respect to a known high-resolution reference image, comparing our results with those of other upsampling methods; second, we test our method for upsampling images of factors of 4 and 8 from their original resolution.

Quantitative performance. For the experiments we use four 512×512 grayscale images: *Lena*, *Barbara*, *Peppers*, and *Gaussian disc*². The images were first downsampled 4 times by 2 stages of wavelet decomposition, obtaining y_{low} ; then the images are upsampled to the original size from y_{low} . Both the one-stage (1) and the m -stage (2) algorithm are tested. We compare them against two other recent upsampling methods: the contourlet-based upsampling

²The *Lena*, *Barbara*, and *Peppers* images are from <http://www.cs.tut.fi/~lasip/2D>, whereas *Gaussian disc* is from [15]. All the experiments in the present paper were carried out using these four images. We note that our first three images differ from the corresponding ones used by the authors of [15].

	\hat{y}_{ℓ^2}	Luong et al. [14]	Contourlet [15]	Proposed	
				one-stage	m -stage
<i>Lena</i>	29.34	28.98	29.79	30.40	30.53
<i>Barbara</i>	23.66	23.54	23.70	23.84	23.86
<i>Peppers</i>	30.18	29.89	30.52	31.74	32.13
<i>Gaus.disc</i>	40.62	41.20	41.70	44.56	49.51

Table 1: PSNR (dB) of the images upsampled 4 times from the approximation coefficients of a two-level wavelet decomposition of the original high-resolution y using a periodically padded Symlet transform of order 8.

	\hat{y}_{ℓ^2}	Luong et al. [14]	Contourlet [15]	Proposed	
				one-stage	m -stage
<i>Lena</i>	0.826	0.819	0.834	0.847	0.849
<i>Barbara</i>	0.664	0.661	0.668	0.683	0.686
<i>Peppers</i>	0.817	0.812	0.820	0.842	0.845
<i>Gaus.disc</i>	0.986	0.986	0.989	0.995	0.998

Table 2: SSIM [17] of the images upsampled 4 times from the approximation coefficients of a two-level wavelet decomposition of the original high-resolution y using a periodically padded Symlet transform of order 8.

algorithm³ [15] and the non-local interpolation algorithm proposed by Luong et al.⁴ [14]. The PSNR and SSIM [17] results are summarized in Tables 1 and 2. In order to avoid the influence of border distortions and thus provide more fair comparisons, PSNR and SSIM values are calculated over the central part of the images, omitting a border of 15-pixel width. Both versions of our algorithm achieve better quantitative results than the ones provided for comparison [14],[15]. In particular, the m -stage version demonstrates the best performance for all images. The numerical difference between the two versions is noticeable especially for the *Gaussian disc*.

Visual appearance. While the SSIM is known to be superior to the PSNR as a measure of subjective image quality, we argue that neither of these two measures is, in general, faithful enough and that direct visual inspection is necessary to assess the quality of the upsampled estimates. Figures 2 and 3 provide a visual comparison between the different methods. Visual examination shows that the images reconstructed with our algorithm are practically free of any kind of ringing artifacts, look sharp, and no significant geometrical distortions can be seen. Nevertheless, some smearing is present in the areas where the fine details have been smoothed in the low-resolution image (e.g., details of the feathers in Figure 3) especially with the one-stage reconstruction. With the m -stage algorithm the smearing is less pronounced, providing visually better images. Despite its relatively high PSNR and SSIM scores, it should be clear from the figures that the \hat{y}_{ℓ^2} estimate is actually the worst among these five.

Implementation and computational complexity. In all our experiments Symlets of order 8 were utilized as the wavelet basis. In order to obtain exactly $n_1 \times n_2$ coefficients for the low-resolution image, periodic padding is used.

A one-step realization of the original BM3D algorithm [3] is used, without the collaborative Wiener filtering step. The filter's internal 3D transform is a combination of a 2D DCT transform applied to each block and of a 1D Haar wavelet transform applied along the third dimension. Other filter parameters were fixed through all experiments, but different for one- and m -stage algorithms. The parameters are as follows.

One-stage algorithm. Experimentally, we found that the algorithm performs better if we start with larger blocks and then switch to a smaller size. Thus, 8×8 block-size was used for the first 7 iter-




Figure 2: Results of reconstruction (upsampling 4 times) of *Lena* (face). Pictures are cropped from upsampled 512×512 size images.

ations, and 5×5 for others. A total of $k_{\text{final}} = 30$ iterations were performed. The σ_k parameter was set to decrease linearly starting from 20 with step $\Delta\sigma = 0.3$.

m -stage algorithm. The filter parameters for each stage are given in Table 3. These parameters ensure both good performance and graceful convergence of the reconstruction. Note that although the size of the block progressively increases with the stages, its relative size with respect to the image $\hat{y}_j^{(k-1)}$ to be filtered is actually decreasing, because the size of $\hat{y}_j^{(k-1)}$ doubles at every stage. The decrease in the (relative) block-size is consistent with the analyses given in [2] and [16].

The BM3D filter is implemented in C++ and compiled into a dll library. The rest of the code is realized in Matlab. Upsampling a 128×128 image to 512×512 size on 2-GHz Intel Core 2 Duo processor takes about 117 seconds with the one-stage algorithm and about 86 seconds with the m -stage algorithm. However, the current software is not optimized to use more than one processing core and the 2D DCT transform inside the BM3D filter is implemented just as matrix multiplications.

Improvement rate. As illustrated in the plot in Figure 4, our experiments show that the one-stage algorithm in only 5 iterations outperforms the algorithms [14] and [15] and that after 15 iterations the

³Matlab implementation available online at

<http://www.ifp.uiuc.edu/~nmueler/interpcontourlet.zip>

⁴Upsampled images have been kindly provided by Hiep Luong.




Figure 3: Results of reconstruction (upsampling 4 times) of *Lena* (hat). Pictures are cropped from upsampled 512×512 size images.

Parameters	stage		
	1	2	3
k_{final}	20	20	20
σ_0	35	25	25
$\Delta\sigma$	0.5	0.3	0.3
block size	3	5	8

Table 3: Parameters used in the m -stage upsampling algorithm.

PSNR is already quite close to its final value. It is worth to comment about the little PSNR improvement achieved for *Barbara*. This happens because downsampling the image four times completely blurs the textures (e.g., the stripes of the trousers), making their reconstruction practically impossible.

Upsampling from original resolution. In the previous experiments we considered upsampling from downsampled images obtained as wavelet approximation coefficients of a given high-resolution image. Let us present also images upsampled of factors 4 and 8 from their original resolution. It should be emphasized that in this case we are not performing interpolation, instead we seek a high-resolution image whose wavelet approximation coefficients in the subband $\Omega_{L^m L^m}$ coincide (up to scaling factor 2^m) to the pixel




Figure 4: PSNR gain (dB) vs. number of iterations for the one-stage algorithm.

values of the given low-resolution image. Figure 5 shows three fragments of the *Cameraman*, *Text*, and *Lighthouse* images at their original resolution. We upsample these fragments applying the m -stage algorithm with symmetric padding of the wavelet transform (thus avoiding the border artifacts which would arise if periodic padding were used). The obtained high-resolution images are shown in Figure 6. Again, the visual quality of the upsampled images is good: edges are sharp and ringing artifacts are absent.

4. DISCUSSION AND CONCLUSIONS

We have introduced a new upsampling algorithm based on Block-Matching and 3D collaborative filtering. This algorithm exploits the ideas of self-learning and sparsity constraints to reconstruct missing wavelet detail coefficients and the idea of the iterative projections on the constraint subspaces (alternated projections).

This work was mostly inspired by our previous contribution [6], where a first attempt was made to apply iterative spatially adaptive filtering to the compressed sensing task. The system (1), in fact, is equivalent to the recursive system used in [6], with the exception of a missing excitation-noise term. While the excitation noise played a key role for image reconstruction from *highly incomplete* data, its influence to the upsampling process is found to be minimal.

The proposed upsampling algorithm can be further extended in a number of ways. The block-matching procedure can be applied to find similar blocks not only in the high-resolution approximation image, but also among blocks in the low-resolution image, thus enabling a sort of interscale matching (like in fractal coding). Further modifications also include changing of the observation model from wavelets to DCT (applied globally or locally) or other transforms. In [4], we present an extension of our upsampling algorithm to the more general problem of image and video super-resolution. We wish also to point the reader to an independent work by Li [12] which is concurrent to our present contribution (we have received a preprint after the submission of the draft of our manuscript). In his work, he proposes an interpolation algorithm based on an iterative BM3D filtering scheme similar ours, with the main difference in the observation model, which is there completely in the pixel domain (the low-resolution image is assumed to be an arbitrarily sampled version of the high-resolution one, without antialias filtering). Also in his case, the BM3D enables high-quality interpolation, competitive to the state-of-the-art.

The main drawback of the iterative algorithm remains its speed, which hinders its use for real-time applications. One possible approach to reduce computational complexity can be switching from BM3D interpolation to a lower-complexity interpolation in smooth areas, similar to the way it is done in the NEDI algorithm [10].




Figure 5: Fragments of the *Cameraman*, *Text*, and *Lighthouse* images.




Figure 6: Upsampling of the three fragments shown in Figure 5 with the m -stage algorithm. From top to bottom: *Cameraman* (4 times, $m = 2$), *Text* (4 times, $m = 2$), *Lighthouse* (8 times, $m = 3$).

REFERENCES

- [1] Buades, A., B. Coll, and J. M. Morel, "A review of image denoising algorithms, with a new one", *Multisc. Model. Simulat.*, vol. 4, no. 2, pp. 490-530, 2005.
- [2] Buades, A., B. Coll, J.M. Morel, C. Sbert, "Non local demosaicing", Preprint CMLA Cachan, no. 2007-15 2007.
- [3] Dabov, K., A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3D transform-domain collaborative filtering", *IEEE Trans. Image Process.*, vol. 16, no. 8, Aug. 2007. <http://www.cs.tut.fi/~foi/GCF-BM3D>
- [4] Danielyan, A., A. Foi, V. Katkovnik, and K. Egiazarian, "Image and video super-resolution via spatially adaptive block-matching filtering", *Proc. 2008 Int. Workshop Local Non-Local Approx. Image Process., LNLA2008*, Lausanne, Switzerland, Aug. 2008.
- [5] Ebrahimi, M., E. R. Vrscay, "Solving the Inverse Problem of Image Zooming using Self-Examples", *Int. Conf. Image Analysis and Recognition ICIAR 2007, Lecture Notes in Computer Science*, vol. 4633, pp. 117-130, Montreal, Canada, 2007.
- [6] Egiazarian, K., A. Foi, and V. Katkovnik, "Compressed Sensing Image Reconstruction via Recursive Spatially Adaptive Filtering", *Proc. IEEE Int. Conf. Image Process., ICIP 2007*, San Antonio (TX), USA, pp. 549-552, Sep. 2007.
- [7] Freeman, W.T., T.R. Jones, and E.C. Pasztor, "Example-based super-resolution", *IEEE Comp. Graphics Applications*, vol. 22, no. 2, pp. 56-65, 2002.
- [8] Jiji, C. V., M. V. Joshi, and S. Chaudhuri, "Single-frame image super-resolution using learned wavelet coefficients", *Int. J. Imaging Syst. Technology*, vol. 14, no. 3, pp. 105-112, 2004.
- [9] Jiji, C. V., and S. Chaudhuri, "Single-frame image super-resolution through contourlet learning", *EURASIP J. Appl. Signal Process.*, Jan. 2006.
- [10] Li, X., and M. T. Orchard, "New edge-directed interpolation", *IEEE Trans. Image Process.*, vol. 10, no. 10, pp. 1521-1527, Oct. 2001.
- [11] Li, X., "Image resolution enhancement via data-driven parametric models in the wavelet space", *J. Image Video Process.*, Jan. 2007.
- [12] Li, X., "Patch-based Nonlocal Image Interpolation", *Proc. 2008 Int. Workshop Local Non-Local Approx. Image Process., LNLA2008*, Lausanne, Switzerland, Aug. 2008.
- [13] Liu, H.C., Y. Feng, and G.Y. Sun, "Wavelet Domain Image Super-Resolution Reconstruction Based on Image Pyramid and Cycle-Spinning", *J. Phys.: Conf. Ser.*, 48, pp. 417-421, 2006
- [14] Luong, H.Q., A. Ledda and W. Philips, "An Image Interpolation Scheme For Repetitive Structures", *Int. Conf. Image Analysis and Recognition, ICIAR 2006, Lecture Notes in Computer Science*, vol. 4141, pp. 104-115, Povo de Varzim, Portugal, Sep. 2006.
- [15] Mueller, N., Y. Lu, and M. N. Do, "Image interpolation using multiscale geometric representations", *Proc. SPIE Electronic Imaging*, San Jose (CA), USA, 2007.
- [16] Protter, M., M. Elad, H. Takeda, and P. Milanfar, "Generalizing the Non-Local-Means to Super-Resolution Reconstruction", to appear in *IEEE Trans. Image Process.*, 2008.
- [17] Wang, Z., A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli, "Image quality assessment: From error visibility to structural similarity", *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600-612, Apr. 2004.