Astronomical Data Analysis Software and Systems XIII
ASP Conference Series, Vol. 314, 2004
F. Ochsenbein, M. Allen, and D. Egret, eds.

Wavelet-Based Superresolution in Astronomy

Rebecca Willett, Robert Nowak

Rice University Department of Electrical and Computer Engineering, Houston, TX and University of Wisconsin Department of Electrical and Computer Engineering, Madison, WI

Ian Jermyn, Josiane Zerubia

INRIA, Project Ariana, Sophia-Antipolis, France

Abstract. High-resolution astronomical images can be reconstructed from several blurred and noisy low-resolution images using a computational process known as superresolution reconstruction. Superresolution reconstruction is closely related to image deconvolution, except that the low-resolution images are not registered and their relative translations and rotations must be estimated in the process. The novelty of our approach to the superresolution problem is the use of wavelets and related multiresolution methods within an expectation-maximization reconstruction process to improve the accuracy and visual quality of the reconstructed image. Simulations demonstrate the effectiveness of the proposed method, including its ability to distinguish between tightly grouped stars with a small set of observations.

1. Introduction

The physical resolution of astronomical imaging devices such as space telescopes is limited by system parameters such as lens aperture and CCD array properties and by physical effects such as the atmosphere and the interstellar/intergalactic medium. However, such systems typically do not take a single snapshot of a celestial body, but rather collect a series of images. Due to mechanical vibrations of the instrument and movement of the satellite relative to the body being images, the images collected are all slightly different observations of the same scene. Superresolution image reconstruction refers to the process of reconstructing a new image with a higher resolution using this collection of low resolution, shifted, rotated, and often noisy observations. This allows users to see image detail and structures which are difficult if not impossible to detect in the raw data.

Superresolution is a useful technique in a variety of applications (Schultz & Stevenson, 1996; Hardie, Barnard, & Armstrong, 1997), and recently, researchers have begun to investigate the use of wavelets for superresolution image reconstruction (Nguyen, Milanfar, & Golub, 2001). We present a new method for superresolution image reconstruction based on the wavelet transform in the presence of Gaussian noise and on an analogous multiscale approach in the presence of Poisson noise. To construct the superresolution image, we use an approach based on maximum penalized likelihood estimation. The reconstructed image is the argument (in our case the superresolution image) that maximizes the sum of a log-likelihood function and a penalizing function. The penalizing function can be specified by an ad hoc smoothness measure, a Bayesian prior distribution for the image (Hebert & Leahy, 1989; Green, 1990), or a complexity measure (Liu & Moulin, 2001). Smoothness measures include simple quadratic functions that measure the similarity between the intensity values of neighboring pixels, as well as non-quadratic measures that better preserve edges. Similar penalty functions result from Markov Random Field (MRF) priors. Complexity measures are usually associated with an expansion of the intensity image with respect to a set of basis functions (e.g. Fourier or wavelet) and count the number of terms retained in a truncated or pruned series (Saito, 1994; Krim & Schick, 1999); the more terms (basis functions) used to represent the image, the higher the complexity measure. Many algorithms (e.g. Expectation-Maximization algorithms, the Richardson Lucy algorithm, or close relatives) have been developed to compute MPLEs under various observation models and penalization schemes.

Wavelets and multiresolution analysis are especially well-suited for astronomical image processing because they are adept at providing accurate, sparse representations of images consisting of smooth regions with isolated abrupt changes or singularities (*e.g.* stars against a dark sky). Many investigators have considered the use of wavelet representations for image denoising, deblurring, and image reconstruction; for examples, see Mallat, 1998, and Starck, Murtagh, & Bijaoui, 1998. The proposed approach uses an EM algorithm for superresolution image reconstruction based on a penalized likelihood formulated in the wavelet domain. Regularization is achieved by promoting a reconstruction with low-complexity, expressed in terms of the wavelet coefficients, taking advantage of the well known sparsity of wavelet representations.

The EM algorithm proposed here extends the work of Nowak & Kolaczyk, 2000, and Figueiredo & Nowak, 2002, which addressed image deconvolution with a method that combines the efficient image representation offered by the discrete wavelet transform (DWT) with the diagonalization of the convolution operator obtained in the Fourier domain. The algorithm alternates between an E-step based on the fast Fourier transform (FFT) and a DWT-based M-step, resulting in an efficient iterative process requiring $O(N \log N)$ operations per iteration, where N is the number of pixels in the superresolution image.

2. Problem Formulation

In the proposed method, each observation, y_k, is modeled as a shifted, rotated, blurred, downsampled, and noisy version of the superresolution $\boldsymbol{x}$. The shift and rotation is caused by the movement of the instrument, and the blur is caused by the point spread function (PSF) of the instrument optics and the integration done by the CCD array. The downsampling (subsampling) operator models the change in resolution between the observations and the desired superresolution image. If the noise can be modeled as additive white Gaussian noise, then we have the observation model

$$y_k = DBS_k\boldsymbol{x} + n_k, \quad k = 1, \ldots, n$$

where D is the downsampling operator, B is the blurring operator, S_k is the shift and rotation operator for the k^{th} observation, and n_k is additive white Gaussian noise with variance σ^2. By collecting the series of observations into one array, $\boldsymbol{y}$, the noise observations into another array, $\boldsymbol{n}$, and letting H be a block diagonal matrix composed of the n matrixes DBS_k for $k = 1, \ldots, n$, then we have the model

$$\boldsymbol{y} = H\boldsymbol{x} + \boldsymbol{n}. \tag{1}$$

Similarly, if the noise is better modeled a Poisson, then we have the model

$$\boldsymbol{y} \sim \text{Poisson}(H\boldsymbol{x}). \tag{2}$$

From the formulation above, it is clear that superresolution image reconstruction is a type of inverse problem in which the operator to be inverted, H, is partially unknown due to the unknown shifts and rotations of the observations. The first step of our approach is to estimate these parameters by registering the low-resolution observations to one another. Using these estimates, we reconstruct an initial superresolution image estimate $\widehat{\boldsymbol{x}}$. This estimate is used in the third step, where we re-estimate the shift and rotation parameters by registering each of the low resolution observations to the initial superresolution estimate. Finally, we use a wavelet-based EM algorithm to solve for $\widehat{\boldsymbol{x}}$ using the registration parameter estimates. We begin by describing a wavelet-based method for the Gaussian noise model, and follow that by a discussion of a multiscale technique for Poisson data. Each of these steps is detailed below.

3. Registration of the Observations

The first step in the proposed method is to register the observed low-resolution images to one another using a Taylor series expansion. This was proposed by Irani and Peleg, 1991. In particular, let f_1 and f_2 be the continuous images underlying the sampled images y_1 and y_2, respectively. If f_2 is equal to a shifted, rotated version of f_1, then we have the relation

$$f_2(t_x, t_y) = f_1(t_x \cos r - t_y \sin r + s_x, t_y \cos r + t_x \sin r + s_y).$$

where (s_x, s_y) is the shift and r is the rotation. A first order Taylor series approximation of f_2 is then

$$\widehat{f_2}(t_x, t_y) = f_1(t_x, t_y) + \left(s_x - t_x r - t_y r^2/2\right)\frac{\partial f_1}{\partial t_x} + \left(s_y - t_y r - t_x r^2/2\right)\frac{\partial f_1}{\partial t_y}.$$

Let $\widehat{y}_2$ be a sampled version of $\widehat{f_2}$; then y_1 and y_2 can be registered by finding the s_x, s_y, and r which minimize $\|y_1 - \widehat{y}_2\|_2^2$, where $\|x\|_2^2 = \sum_i x_i^2$. This minimization is calculated with an iterative procedure which ensures that the motion being estimated at each iteration is small enough for the Taylor series approximation to be accurate; see (Irani & Peleg, 1991) for details. This method was applied

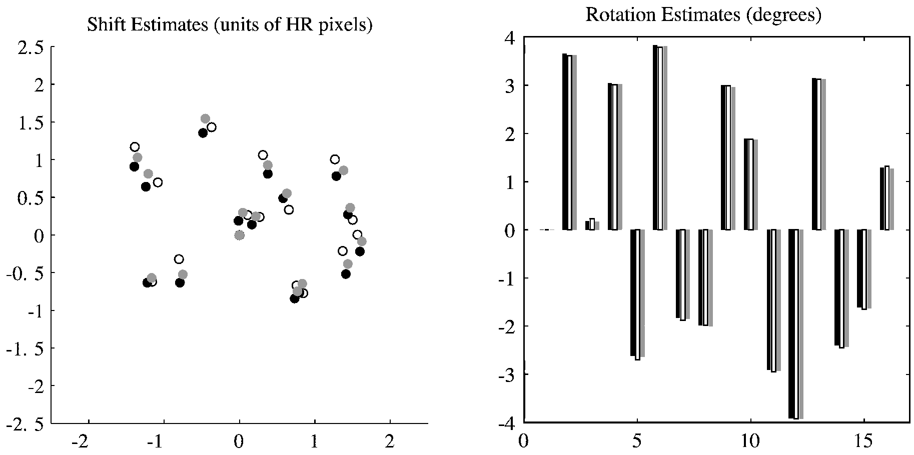

(a) (b)

Figure 1. Image registration results. (a) True shifts (black), initial shift estimates (hollow), and final shift estimates (gray). (b) True rotations (black), initial rotation estimates (hollow), and final rotation estimates (gray).

to the earth image data in Figure 2. The sixteen true shifts and rotations are displayed in black in Figure 1, and the results of this registration method are displayed in hollow circles and bars.

After the registration parameters have been initially estimated using the above method, we use these estimates to calculate an initial superresolution image as described in Section 4.. This initial image estimate is then used to refine the registration parameter estimates. The method is the same as above, but instead of registering a low resolution estimate, y_2, to another low resolution estimate, y_1, we instead register it to $DBS_1\widehat{\boldsymbol{x}}$. The results of this refinement are displayed in gray in Figure 1. From these plots, it is clear that the Taylor series based approach can produce highly accurate results. However, in low SNR scenarios, where confidence in registration parameter estimates may be low, the estimates can be further refined at each iteration of the proposed EM algorithm, as discussed in the following sections.

Note that the motion model considered here encompasses only shift and rotation movement. When recording still or relatively still objects distant from the imaging device, this model is sufficient. More sophisticated models are an area of open research.

4. Multiscale Expectation-Maximization

Maximization is facilitated within the EM framework through the introduction of a particular "unobservable" or "missing" data space. The key idea in the EM algorithm is that the indirect (inverse) problem can be broken into two subproblems; one which involves computing the expectation of unobservable data (as though no blurring or downsampling took place) and one which entails estimating the underlying image from this expectation. By carefully defining the unobservable data for the superresolution problem, we derive an EM algorithm which consists of linear filtering in the E-step and image denoising in the M-step.

The Gaussian observation model in (1) can be written with respect to the DWT coefficients $\boldsymbol{\theta}$, where $\boldsymbol{x} = W\boldsymbol{\theta}$ and W denotes the inverse DWT operator (Mallat, 1998):

$$\boldsymbol{y} = HW\boldsymbol{\theta} + \boldsymbol{n}.$$

Clearly, if we had $\boldsymbol{y} = W\boldsymbol{\theta}+\boldsymbol{n}$ (*i.e.* if no subsampling or blurring had occurred), we would have a pure image denoising problem with white noise, for which wavelet-based denoising techniques are very fast and nearly optimal (Mallat, 1998). Next note that the noise in the observation model can be decomposed into two different Gaussian noises (one of which is non-white):

$$\boldsymbol{n} = \alpha H\boldsymbol{n}_1 + \boldsymbol{n}_2$$

where α is a positive parameter, and $\boldsymbol{n}_1$ and $\boldsymbol{n}_2$ are independent zero-mean Gaussian noises with covariances $\Sigma_1 = I$ and $\Sigma_2 = \sigma^2 I - \alpha^2 HH^T$, respectively. Using $\boldsymbol{n}_1$ and $\boldsymbol{n}_2$, we can rewrite the Gaussian observation model as

$$\boldsymbol{y} = H\underbrace{(W\boldsymbol{\theta} + \alpha\boldsymbol{n}_1)}_{\boldsymbol{z}} + \boldsymbol{n}_2.$$

This observation is the key to our approach since it suggests treating $\boldsymbol{z}$ as missing data and using the EM algorithm to estimate $\boldsymbol{\theta}$.

An analogous formulation is possible for the Poisson noise model. In this case, photon projections can be described statistically as follows. Photons are emitted (from the emission space) according to a high resolution intensity $\boldsymbol{x}$. Those photons emitted from location m are detected (in the detection space) at position n with transition probability $H_{m,n}$, where H is the superresolution operator from (2). In such cases, the measured data are distributed according to

$$\boldsymbol{y}_n \sim \text{Poisson}\left(\sum_m H_{m,n}x_m\right). \tag{3}$$

In this formulation of the EM algorithm, the missing data is defined as $\boldsymbol{z} = \{z_{m,n}\}$, where $z_{m,n}$ denotes the number of photons emitted from m and detected at n. The complete data are Poisson distributed according to

$$z_{m,n} \sim \text{Poisson}(H_{m,n}\boldsymbol{x}_m).$$

Hence the observed data $\boldsymbol{y}$ in (3) are given by $\boldsymbol{y}_n = \sum_m z_{m,n}$. Additionally, were we able to observe $\boldsymbol{z} = \{z_{m,n}\}$, the direct emission data for each location m is given by sums of the form $\sum_n z_{m,n} \sim \text{Poisson}(\boldsymbol{x}_m)$. Therefore, if $\boldsymbol{z}$ were known, we could avoid the inverse problem altogether and simply deal with the issue of estimating a Poisson intensity given direct observations.

From these formulations of the problem for Gaussian and Poisson data, the EM algorithm produces a sequence of estimates $\{\boldsymbol{x}^{(t)}, t = 0, 1, 2, \ldots\}$ by alternately applying two steps:

E-Step: Updates the estimate of the missing data using the relation:

$$\widehat{\boldsymbol{z}}^{(t)} = E\left[\boldsymbol{z}|\boldsymbol{y}, \widehat{\boldsymbol{\theta}}^{(t)}\right].$$

In the Gaussian case this can be reduced to a Landweber iteration (Landweber, 1951):

$$\widehat{\boldsymbol{z}}^{(t)} = \widehat{\boldsymbol{x}}^{(t)} + \frac{\alpha^2}{\sigma^2} H^T \left(\boldsymbol{y} - H\widehat{\boldsymbol{x}}^{(t)}\right).$$

Here, computing $\widehat{\boldsymbol{z}}^{(t)}$ simply involves applications of the operator H. Recall that H consists of shifting, blurring, and rotation (which can be computed rapidly with the 2D FFT) and downsampling (which can be computed rapidly in the spatial domain). Thus the complexity of each E-Step is $O(N \log N)$.

In the Poisson case, this step is reduced to

$$z_{n,m}^{(t)} = \frac{\boldsymbol{y}_n \widehat{\boldsymbol{x}}_m^{(t)} H_{n,m}}{\sum_\ell \widehat{\boldsymbol{x}}_\ell^{(t)} H_{n,\ell}},$$

which can also be computed in $O(N \log N)$ operations.

M-Step: Updates the estimate of the superresolution image $\boldsymbol{x}$. In the Gaussian case, this constitutes updating the wavelet coefficient vector $\boldsymbol{\theta}$ according to

$$\widehat{\boldsymbol{\theta}}^{(t+1)} = \arg\min_{\boldsymbol{\theta}} \left\{ \frac{\|W\boldsymbol{\theta} - \widehat{\boldsymbol{z}}^{(t)}\|_2^2}{2\alpha^2} + \mathrm{pen}(\boldsymbol{\theta}) \right\}$$

and setting $\widehat{\boldsymbol{x}}^{(t+1)} = W\widehat{\boldsymbol{\theta}}^{(t+1)}$. This optimization can be preformed using any wavelet-based denoising procedure. For example, under an *i.i.d.* Laplacian prior, $\mathrm{pen}(\boldsymbol{\theta}) = -\log p(\boldsymbol{\theta}) \propto \tau\|\boldsymbol{\theta}\|_1$ (where $\|\boldsymbol{\theta}\|_1 = \sum_i |\boldsymbol{\theta}_i|$ denotes the l_1 norm), $\widehat{\boldsymbol{\theta}}^{(t+1)}$ is obtained by applying a *soft-threshold* function to the wavelet coefficients of $\widehat{\boldsymbol{z}}^{(t)}$. For the simulations presented in this paper, we applied a similar denoising method described in (Figueiredo & Nowak, 2001), which requires $O(N)$ operations.

In the Poisson case, the M-Step is equivalent to photon-limited image denoising. In practice, this commonly consists of performing maximum likelihood estimation, but these estimates are known to diverge from the true image after several iterations. The denoising can also be accomplished using the Haar-based maximum penalized likelihood estimation method we developed in Willett & Nowak, 2003; the MPLE function employed here is

$$L(\boldsymbol{x}) \;\equiv\; \log p(\boldsymbol{y} \,|\, \sum_m z_{m,n}^{(t)}) \;-\; \frac{1}{2}\,\mathrm{pen}(\sum_m z_{m,n}^{(t)}) \log n, \qquad (4)$$

where $p(\boldsymbol{y}\,|\,\boldsymbol{x})$ is the Poisson likelihood of observing photon counts $\boldsymbol{y}$ given intensities $\sum_m z_{m,n}^{(t)}$ and $\mathrm{pen}(\sum_m z_{m,n}^{(t)})$ is the number of non-zero coefficients in the Haar-based multiscale representation of $\sum_m z_{m,n}^{(t)}$ and n is the total number of photons in the observation vector $\boldsymbol{y}$. In maximizing this function, the resulting reconstruction will be one that has a relatively high likelihood value as well as a relatively low complexity Haar representation. This denoising procedure is a bottom-up tree pruning algorithm which requires $O(N \log N)$ operations.

In both the Gaussian and Poisson noise cases, the proposed method has the advantage that the M-step is a denoising procedure, and the multiscale methods employed here are both near-minimax optimal.

In low SNR cases, where confidence in the initial estimates of the shift and rotation parameters may be low, the proposed algorithm can be modified by simply inserting an additional step in which the parameter estimates are updated based on the current estimate of the superresolution image $\widehat{\boldsymbol{x}}^{(t)}$, as described in the previous section. Similarly, if the blurring operator is only partially known, its parameters can also be updated at each iteration of the proposed algorithm. The resulting scheme is not a standard EM algorithm, but it is guaranteed not to decrease the penalized likelihood function.

5. Simulation Results

We reconstructed two superresolution images to demonstrate the practical effectiveness of the proposed algorithms. A sample low-resolution satellite image of earth is displayed in Figure 2(a). Sixteen such 64×64 images were generated using an original 256×256 image (Figure 2(c), with values ranging from 0 to 255), which was shifted and rotated by a random, subpixel amount, distorted by a 4×4 uniform blur, and contaminated with additive white Gaussian noise with variance $\sigma^2 = 1/2$. The superresolution image in Figure 2(b) was reconstructed using the wavelet-based EM algorithm described above for $\alpha^2 = \sigma^2 = 1/2$. In the simulation, the estimate is initialized with a least-squares superresolution estimate of relatively poor quality. While not presented here, experiments have shown that the proposed approach is competitive with the state of the art in superresolution image reconstruction.

In our second simulation, we studied the effect of the proposed method on an image of stars. The original image is displayed in Figure 3(c). The data for this simulation was generated using the same procedure as for the Earth image. Note from the sample observation image in Figure 3(a) that several stars are indistinguishable prior to superresolution image reconstruction, but that after the application of the proposed method these stars are clearly visible in Figure 3(b). The superresolution image in Figure 3(d) was generated using the same procedure but with only two observation images. Clearly the quality of this result is less than the quality achievable with more observations, but significantly more stars are distinguishable in the output than in any one of the observations. This implies that the proposed method may be useful even for small sets of observations. Finally, Figure 3(e) displayed the output of our multiscale method for Poisson noise. In the case, the mean photon count was 527. As in the Gaussian case, the improvement in resolution is significant.

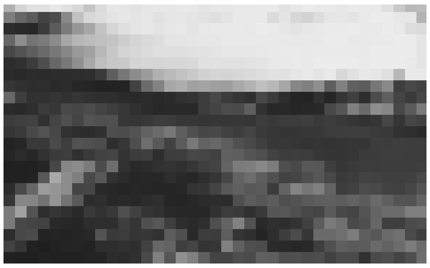

(a)

(b)

(c)

Figure 2. Superresolution results for earth image. (a) One of 16 observation images (64×64), contaminated with Gaussian noise, $\sigma^2 = 0.5$ and a 4×4 uniform blur. (b) 256×256 result. (c) True high resolution image.

6. Conclusions and Future Work

In this paper we present a multiscale approach for superresolution image reconstruction in the presence of either Gaussian or Poisson noise. This approach uses several shifted, rotated, blurred, noisy observations to construct a new image with higher resolution than any of the observations. Because of the multiscale complexity regularization used in the the EM reconstruction algorithm, our approach is robust to noise. This is demonstrated for several practical examples. Notably, we demonstrated that stars which are irresolvable in any individual observation can be clearly distinguished after superresolution image reconstruction, even when using only two observation images.

Future work includes speeding the convergence of the proposed method using a new technique described in (Salakhutdinov & Roweis, 2003) and a more Bayesian approach in which the shift and rotation parameters are integrated out of the problem formulation. In addition, there are several questions still to be addressed, including: How does the blur radius impact reconstruction quality? How much can resolution be recovered at what accuracy? How can the proposed multiscale methods be extended to optimally process data collected by photon detector cells with spatially varying sensitivities? Despite these open areas of research, however, we feel that wavelet-based superresolution has the potential to make a significant impact on astronomical imaging.

References

Donoho, D., 1999, Ann. Statist., 27, 859

Figueiredo, M. & Nowak, R., 2002, IEEE Transactions on Image Processing, 29, 1033

Figueiredo, M. & Nowak, R., 2001, IEEE Transactions on Image Processing, 10, 1322

Green, P. J., 1990, IEEE Trans. Med. Imaging, 9, 84

Hardie, R. C., Barnard, K. J., & Armstrong, E. E., 1997, IEEE Transactions on Image Processing, 6, 1621

Hebert, T. & Leahy, R., 1989, IEEE Trans. Med. Imaging, 8, 194

Irani, M. & Peleg, S. 1991, Computer Vis. Graph. Image Process.: Graph. Models Image Process, 53, 231

Kolaczyk, E., 1999, J. Amer. Statist. Assoc., 94, 920

Kolaczyk, E. & Nowak, R., 2003, to appear in the Annals of Statistics, "Multiscale Likelihood Analysis and Complexity Penalized Estimation", `http://www.ece.wisc.edu/~nowak/pubs.html`

Krim, H., & Schick, I. C., 1999, IEEE Trans. on Information Theory, 45

Landweber, L., 1951, American Journal of Mathematics, 73, 615

Liu, J. & Moulin, P., 2001, IEEE Transactions on Image Processing, 10, 841

Mallat, S., 1998, A Wavelet Tour of Signal Processing

Nowak, R., & Kolaczyk, E., 2000, IEEE Transactions on Information Theory, 46, 1811

Nguyen, N., Milanfar, P., & Golub, G., 2001, IEEE Transactions on Image Processing, 10, 573

Saito, N, 1994, Wavelets in Geophysics, ed. Foufoula-Georgiou and Kumar

Salakhutdinov, R. & Roweis, S., 2003, International Conference on Machine Learning, 664

Schultz, R. R. & Stevenson, R. L., 1996, IEEE Transactions on Image Processing, 5, 996

Starck, J., Murtagh, F., & Bijaoui, A., 1998, Image Processing and Data Analysis: The Multiscale Approach

Timmermann, K. & Nowak, R, 1999, IEEE Transactions on Information Theory, 45, 846

Willett, R. & Nowak, R., 2003, IEEE Transactions of Medical Imaging, 22, 332

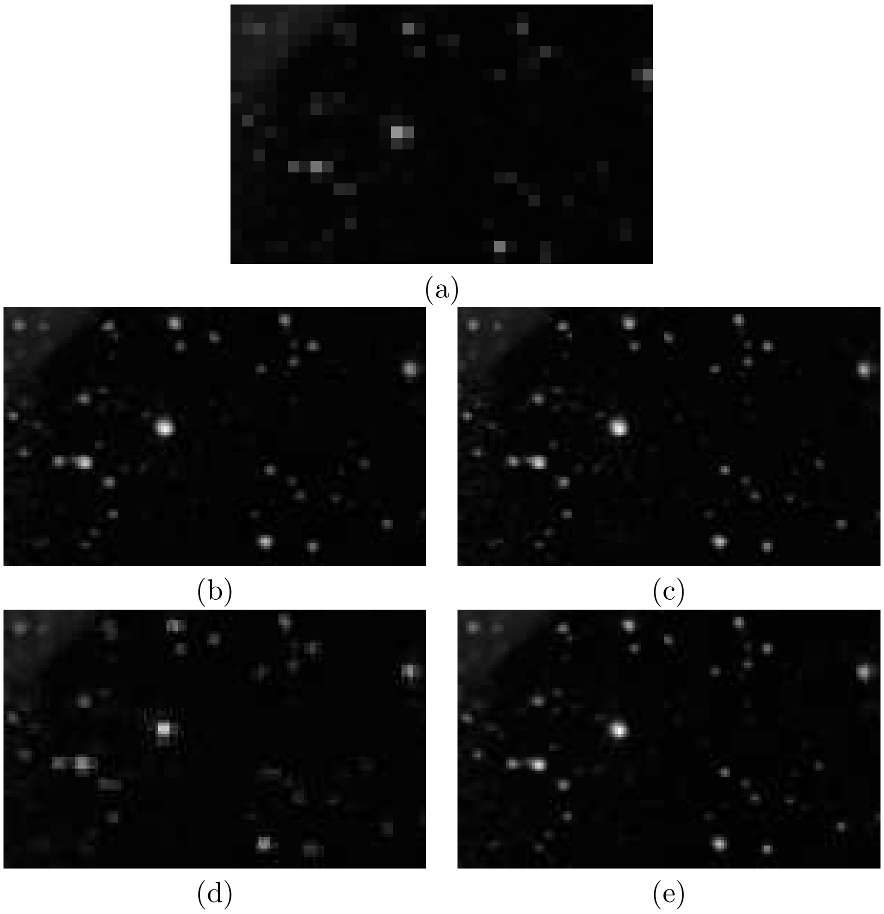

Figure 3. Superresolution results for stars image. (a) One of 16 observation images (64×64), contaminated with Gaussian noise, $\sigma^2 = 0.5$ and a 4×4 uniform blur. (b) 256×256 result. (c) True high resolution image. (d) 256×256 result obtained using only two 64×64 observation images. (e) 256×256 result obtained using 16 64×64 observation images contaminated with Poisson noise; mean photon count per observation image = 527.