4th AIP International Conference and the 1st Congress of the IPIA IOP Publishing
Journal of Physics: Conference Series **124** (2008) 012037 doi:10.1088/1742-6596/124/1/012037

Robust Super-Resolution by Minimizing a Gaussian-weighted L_2 Error Norm

Tuan Q. Pham[1], **Lucas J. van Vliet**[2], **Klamer Schutte**[3]
[1] *Canon Information Systems Research Australia, 1 Thomas Holt drive, North Ryde, NSW 2113, Australia*
[2] *Quantitative Imaging Group, Department of Imaging Science and Technology, Faculty of Applied Sciences, Delft University of Technology, Lorentzweg 1, 2628 CJ Delft, The Netherlands*
[3] *Electro-Optics Group, TNO Defence, Security and Safety, P.O. Box 96864, 2509 JG The Hague, The Netherlands*

Abstract. Super-resolution restoration is the problem of restoring a high-resolution scene from multiple degraded low-resolution images under motion. Due to imaging blur and noise, this problem is ill-posed. Additional constraints such as smoothness of the solution via regularization is often required to obtain a stable solution. While adding a regularization term to the cost function is a standard practice in image restoration, we propose a restoration algorithm that does not require this extra regularization term. The robustness of the algorithm is achieved by a Gaussian-weighted L_2-norm in the data misfit term that does not response to intensity outliers. With the outliers suppressed, our solution behaves similarly to a maximum-likelihood solution in the presence of Gaussian noise. The effectiveness of our algorithm is demonstrated with super-resolution restoration of real infrared image sequences under severe aliasing and intensity outliers.

1. Introduction

Image restoration belongs to the class of inverse problems. It aims to reconstruct the real underlying distribution of a physical quantity called the *scene* from a set of measurements. In Super-Resolution (SR) restoration [16], the scene is restored from a series of degraded Low-Resolution (LR) images, each of which is a projection of the real and continuous scene onto an array of sensor elements. Although the original scene is continuous, image restoration only produces a digital image representation. The algorithm uses a forward model which describes how the LR observations are obtained from the unknown High-Resolution (HR) scene. This model often incorporates the Point Spread Function (PSF) of the optics, the finite-size photo-sensitive area of each pixel sensor, and the stochastic noise sources.

Two types of degradations are present in a LR image. The first degradation is the reduction of spatial resolution. This is a collective result of multiple blur sources such as the PSF of the optics, defocusing, motion blur, and sensor integration. The second type of degradation is noise. Due to these degradations, SR is an ill-posed problem which means that the solution is highly sensitive to noise in the observations. A stable solution can only be reached by numerical approximation which may enforce additional constraints such as smoothness of solution [22], a.k.a. regularization.

In the first part of this paper, we present a new solution for robust SR restoration. Different from the common practice of adding a regularization term to the cost function to suppress both Gaussian noise and outliers, we handle each of these intensity variations separately. The influence of outlier observations is suppressed by giving these samples a low contribution weight in subsequent analysis. This tonal-based weighting scheme is similar to the ones used by bilateral filtering [23]. After the outliers are suppressed, our solution behaves in similar way to a maximum-likelihood solution in the presence of additive Gaussian noise only. The iterative restoration is terminated before noise overwhelms the signal. The moderately amplified noise is then removed by a fast signal-preserving filter during a postprocessing step.

The second part of this paper is devoted to an objective evaluation of different SR algorithms. The evaluation relies on image features often found in a HR image such as edges and small blobs. Different SR results are compared on five quantitative measures including noise, edge sharpness, edge jaggedness, blob size, and blob contrast. These measurements are plotted on a radar chart per algorithm to facilitate the comparison of multiple algorithms. Although it is unlikely that a single algorithm outperforms others on all performance measures, the size of the polygon that connects the radar points generally indicates which algorithm performs better on a given dataset.

2. Super-resolution restoration

Image formation is a process in which a continuous input signal, i.e. the scene, is degraded by several blur and noise sources to produce an output image. Although each component along the imaging chain introduces its own degradation, the overall process can be modeled as a convolution of the input signal with a Point Spread Function (PSF) followed by noise addition:

$$y(\mathbf{x}) = w(\mathbf{x}) * f(\mathbf{x}) + n(\mathbf{x}) \tag{1}$$

where $*$ denotes the convolution operation. y, w, f, and n are the measured signal, the PSF, the original signal, and noise, respectively. $\mathbf{x}$ is the coordinate vector of the continuous spatial domain.

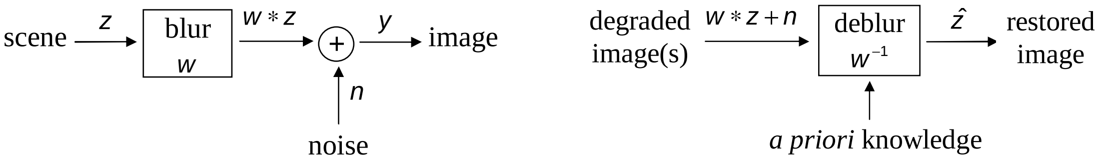

(a) Image formation by blur and additive noise (b) Image restoration by deconvolution

Figure 1. Image formation and restoration models.

Given the signal degradation model (1), image restoration aims to recover an estimate of the scene f from the corrupted image y. To facilitate the use of numerical analysis, all continuous signals are represented by its sampled versions. Because of its infinite bandwidth, however, the continuous scene f is not fully reconstructible. We can only aim to reconstruct a band-limited version of f, of which the sampled version is z (figure 1a). The restored image $\hat{z}$ in figure 1b is then expected to resolve finer details and possess a better Signal-to-Noise Ratio (SNR) than the observation y.

Under a discrete-to-discrete formulation, the imaging model in equation (1) can be rewritten as a matrix multiplication followed by a vector addition of the lexicographically ordered samples:

$$\begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_M \end{bmatrix} = \begin{bmatrix} w_{1,1} & w_{1,2} & \cdots & w_{1,N} \\ w_{2,1} & w_{2,2} & \cdots & w_{2,N} \\ \vdots & \vdots & \ddots & \vdots \\ w_{M,1} & w_{M,2} & \cdots & w_{M,N} \end{bmatrix} \times \begin{bmatrix} z_1 \\ z_2 \\ \vdots \\ z_N \end{bmatrix} + \begin{bmatrix} n_1 \\ n_2 \\ \vdots \\ n_M \end{bmatrix} \tag{2}$$

where the measured value of an m-th sample y_m is a weighted sum of *all* input samples z_r $(r = 1..N)$ plus a noise term:

$$y_m = \sum_{r=1}^{N} w_{m,r} z_r + n_m \qquad \forall\, m = 1..M \tag{3}$$

The weights $w_{m,r}$ are sampled values of the PSF centered at the position of the observed sample y_m. Although this PSF is sampled at all "grid" positions of the HR image z, most of its entries are zero if the PSF has a limited spatial support. The matrix $[w_{i,j}]$ is therefore often sparse. Using bold letters to represent the corresponding matrices, (2) can be written as:

$$\mathbf{y} = \mathbf{W}\mathbf{z} + \mathbf{n} \tag{4}$$

Equation (4) also models the generation of multiple LR images from a single HR scene [8]. In the case of p LR images, $\mathbf{y}$ is a $pM \times 1$ vector of all observed samples from p LR images, $\mathbf{z}$ is still a $N \times 1$ matrix of the desired HR samples, the $pM \times N$ matrix $\mathbf{W}$ represents not only a blur operation but also a geometric transformation and subsampling [3]. The reconstruction of the HR image $\mathbf{z}$ from multiple LR frames $\mathbf{y}$ can therefore be seen as a multi-frame deconvolution operation. In fact, the only difference between multi-frame SR and single-frame deconvolution is the geometric transformation of the LR frames that often leads to irregular sample positions of $\mathbf{y}$. As a result, all derivations for SR reconstruction in this section are also applicable to deconvolution.

If noise is not present, (4) reduces to a linear system which can be solved by singular value decomposition [19] (in a least-squares sense if $pM \neq N$ or $\text{rank}(\mathbf{W}) < N$). However, due to noise and the sizable dimensions of $\mathbf{W}$, numerical approximation is often used to solve $\mathbf{z}$. In the following subsections, we review a maximum-likelihood solution to SR and propose the use of a robust error norm to improve the solution's responses to outliers. Using total variation denoise as a postprocessing step, the new SR algorithm outperforms current methods in handling outliers and noise from input observations.

2.1. Maximum-likelihood super-resolution

Maximum-Likelihood (ML) super-resolution seeks to maximize the probability that the measured images $\mathbf{y}$ are the results of image formation from a common scene $\mathbf{z}$. With the knowledge of the PSF, the simulated LR images can be constructed from $\mathbf{z}$ as $\hat{\mathbf{y}} = \mathbf{W}\mathbf{z}$, this differs from the measured images $\mathbf{y}$ by a noise term only. If the measurement noise is normally distributed with zero mean and variance σ_n^2, the probability that $\mathbf{y}$ was actually measured from $\mathbf{z}$ is:

$$\Pr(\mathbf{y}|\mathbf{z}) = \prod_{m=1}^{pM} \frac{1}{\sigma_n\sqrt{2\pi}} \exp\left(-\frac{(y_m - \hat{y}_m)^2}{2\sigma_n^2}\right) \tag{5}$$

Maximizing (5) is equivalent to minimizing its negative log-likelihood:

$$C(\mathbf{z}) = \frac{1}{2}\sum_{m=1}^{pM} (y_m - \hat{y}_m)^2 = \frac{1}{2}\sum_{m=1}^{pM} e_m^2 \tag{6}$$

where $e_m = y_m - \hat{y}_m = y_m - \sum_{r=1}^{N} w_{m,r} z_r$ is the simulation error at the m-th output pixel. The cost function (6) can also be seen a quadratic error norm between the simulated images $\hat{\mathbf{y}}$ and the measured images $\mathbf{y}$. An ML solution given additive Gaussian distributed noise is therefore defined as:

$$\hat{\mathbf{z}}_{ML} = \arg\min_{\mathbf{z}} C(\mathbf{z}) \tag{7}$$

Equation (7) can be solved by a gradient descent method [12], which starts with an arbitrary initialization $\hat{\mathbf{z}}^0$ and iteratively updates the current estimate in the gradient direction of the cost function:

$$\hat{\mathbf{z}}^{n+1} = \hat{\mathbf{z}}^n - \varepsilon^n \mathbf{g}^n \tag{8}$$

where ε^n is the step size at iteration n, $\mathbf{g}^n = [g_1(\hat{\mathbf{z}}^n)\ g_2(\hat{\mathbf{z}}^n)\ ...\ g_N(\hat{\mathbf{z}}^n)]^T$ is the gradient vector of $C(\mathbf{z}^n)$ whose k-th element is:

$$g_k(\mathbf{z}) = \frac{\partial C(\mathbf{z})}{\partial z_k} = -\sum_{m=1}^{pM} w_{m,k} e_m \qquad \text{for} \quad k = 1..N \tag{9}$$

The ML solution was first applied to SR by Irani and Peleg [9]. It is also known as the iterative back-projection method because the data misfit error at each LR pixel is back-projected to the current HR image to *blame* those HR pixels that contribute to the error. In equation (9), $\hat{\mathbf{z}}^n$ is updated with a convolution of the data misfit error with a back-projection kernel, which is often the PSF itself. If Newton's method [12] is used instead of the gradient descent method, the back-projection kernel is proportional to the square of the PSF. An ML algorithm often converges to a noisy solution. To prevent noise corruption, the iterative process can be stopped as soon as the relative norm of the back-projected error image exceed a certain threshold [8]. This usually takes fewer than five iterations for all examples in this paper.

2.2. Robust maximum-likelihood super-resolution

Because of the unregularized quadratic error norm in the cost function (6), the original ML algorithm is not robust against outliers. As a consequence, ringing and overshoot are often seen in the reconstructed HR image, especially when noise and outliers are present. To reduce such artifacts, many modifications such as regularization [8, 5] and median back-projection [25] have been added to the basic ML solution. Irani [9] also suggested a heuristic approach by removing the contributions of extreme values before back-projecting the errors. In this subsection, we provide a rigorous mathematical basis to support Irani's heuristics. Our SR algorithm minimizes a Gaussian-weighted L_2-norm of the data misfit errors. This robust cost function is minimized by a conjugate-gradient optimization algorithm [12], which results in a much faster convergence than the gradient descent algorithm.

2.2.1. Robust estimation using Gaussian error norm To determine whether an estimator is robust against outliers or not, one needs to review its cost function. The ML solution for SR aims to minimize a quadratic norm of the simulation error $\sum_{m=1}^{pM} \rho_2(e_m)$, where:

$$\rho_2(x) = \tfrac{1}{2}x^2 \qquad\qquad \psi_2(x) = \rho_2'(x) = x \tag{10}$$

The influence function $\psi = \rho'$ defines a "force" that a particular measurement imposes on the solution [7]. For the quadratic norm, the influence function is linearly proportional to the error. A single intensity outlier can therefore bias the solution to an erroneous value. A robust estimator avoids this by reducing the influence of extreme errors to zero. The Gaussian error norm [1], for example, satisfies this requirement:

$$\rho_G(x) = \sigma_t^2 \left(1 - \exp\left(-x^2/2\sigma_t^2\right)\right) \qquad \psi_G(x) = x.\exp\left(-x^2/2\sigma_t^2\right) \tag{11}$$

where the tonal scale σ_t controls the acceptance range of inlier variations. As can be seen from figure 2, the Gaussian norm with $\sigma_t = 1$ resembles the quadratic norm for $x \leq \sigma_t$ and it asymptotically approaches one for $x \geq 3\sigma_t$. The Gaussian influence function is negligible for $x \geq 3\sigma_t$, which means that it is not susceptible to outliers. Apart from its robustness, the Gaussian error norm is also popular for its simplicity. The Gaussian influence function differs from the L_2-norm influence function by only an exponential term (ψ_G in equation (11) compared to ψ_2 (10)). As a consequence, a robust estimator that uses the Gaussian error norm yields a formulation that is comparable to the traditional ML estimator.

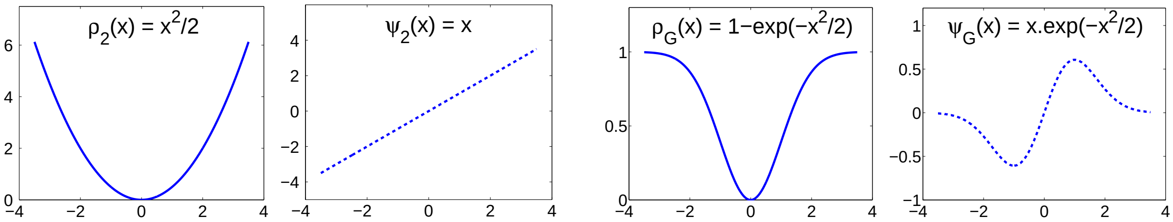

(a) L_2 norm (b) L_2 influence (c) Gaussian norm (d) Gaussian influence

Figure 2. Quadratic error norm versus Gaussian error norm ($\sigma_t = 1$).

Using the Gaussian norm instead of the L_2 norm, the cost function of the robust SR estimator is:

$$C(\mathbf{z}) = \sum_{m=1}^{pM} \rho_G\left(e_m\right) = \sigma_t^2 \sum_{m=1}^{pM} \left[1 - \exp\left(-\frac{e_m^2}{2\sigma_t^2}\right)\right] \tag{12}$$

Its partial derivative w.r.t. an output pixel z_k ($k = 1..N$) is:

$$g_k(\mathbf{z}) = \frac{\partial C(\mathbf{z})}{\partial z_k} = -\sum_{m=1}^{pM} w_{m,k} e_m \exp\left(-\frac{e_m^2}{2\sigma_t^2}\right) = -\sum_{m=1}^{pM} w_{m,k} e_m c_m \tag{13}$$

where $c_m = \exp\left(-e_m^2/2\sigma_t^2\right)$ can be interpreted as a certainty value of the measured sample y_m. This certainty value is close to one for $|e_m| \leq \sigma_t$ and it asymptotically approaches zero for $|e_m| \geq 3\sigma_t$. As a consequence, the contributions of extreme intensity values are suppressed. With this robust weighting scheme, our robust ML solution automatically achieves the heuristic noise reduction scheme suggested by Irani.

2.2.2. Conjugate-gradient optimization Although the robust cost function (12) can be minimized by a gradient descent process in the same way as presented in the previous subsection, a conjugate-gradient optimization algorithm usually converges much faster [12]. The procedure for conjugate-gradient optimization of the cost function (12) is described as follows. Starting with a 3 × 3 median filtered image of a shift-and-add super-fusion result [4], the current HR estimate is updated by:

$$\hat{\mathbf{z}}^{n+1} = \hat{\mathbf{z}}^n + \varepsilon^n \mathbf{d}^n \qquad \text{for} \quad n = 0, 1, 2, \ldots \text{ iterations} \tag{14}$$

where $\mathbf{d}^n = [d_1(\hat{\mathbf{z}}^n)\ d_2(\hat{\mathbf{z}}^n)\ ...\ d_N(\hat{\mathbf{z}}^n)]^T$ is a conjugate-gradient vector, which is derivable from $\mathbf{g}^n$.

One important characteristics of the conjugate-gradient optimization process is that it never updates the HR estimate in the same direction twice. To ensure this condition, the first conjugate-gradient vector is initialized with $\hat{\mathbf{d}}^0 = -\hat{\mathbf{g}}^0$ and all subsequent vectors are computed by the Fletcher-Reeves method [6]:

$$\mathbf{d}^{n+1} = -\mathbf{g}^{n+1} + \beta^n \mathbf{d}^n \qquad \text{where} \qquad \beta^n = \frac{(\mathbf{g}^{n+1})^T \mathbf{g}^{n+1}}{(\mathbf{g}^n)^T \mathbf{g}^n} \tag{15}$$

The optimal step size ε^n at the n-th iteration is calculated by minimizing the next iteration's cost function $C(\hat{\mathbf{z}}^{n+1})$ with respect to ε^n. Setting the derivative of the cost function w.r.t the step size to zero and using differentiation by parts, we obtain:

$$0 = \frac{\partial C(\hat{\mathbf{z}}^{n+1})}{\partial \varepsilon^n} = \sum_{k=1}^{N} \frac{\partial C(\hat{\mathbf{z}}^{n+1})}{\partial \hat{z}_k^{n+1}} \frac{\partial \hat{z}_k^{n+1}}{\partial \varepsilon^n} = \sum_{k=1}^{N} g_k\left(\hat{\mathbf{z}}^{n+1}\right) d_k\left(\hat{\mathbf{z}}^n\right) \tag{16}$$

Equation (16) confirms that the new gradient-descent direction $\mathbf{g}^{n+1}$ is orthogonal to the last update direction $\mathbf{d}^n$. After some further expansion of (16) [8], the optimal step size is found to be:

$$\varepsilon^n = -\sum_{m=1}^{pM} \gamma_m e_m^n c_m^{n+1} \Bigg/ \sum_{m=1}^{pM} \gamma_m^2 c_m^{n+1} \tag{17}$$

where $\gamma_m = \sum\limits_{r=1}^{N} w_{m,r} g_m(\hat{\mathbf{z}}^n)$ is a convolution of the gradient image at iteration n with the PSF. Although ε^n depends on a *future* certainty image $\mathbf{c}^{n+1} = [c_1^{n+1}\ c_2^{n+1}\ ...\ c_{pM}^{n+1}]^T$, $\mathbf{c}^{n+1}$ can be approximated by $\mathbf{c}^n$ because the sample certainty remains stable (especially at later iterations).

The iterative optimization process is stopped after a maximum number of iterations is reached or when the relative norm of the update image against the output image is small enough: $\xi^n = |\hat{\mathbf{z}}^{n+1} - \hat{\mathbf{z}}^n| / |\hat{\mathbf{z}}^n| < T$ (e.g., $T = 10^{-4}$). Since we use an approximation of the future confidence $\mathbf{c}^{n+1}$, the relative norm ξ may not always reduce. To prevent a possible divergence of the robust ML solution, we also terminate the iterative process when the relative norm ξ increases. A summary of the conjugate-gradient optimization algorithm is given in table 1. Note that all formula from (13) to (15) of the robust ML solution differ from the corresponding formula of the traditional ML solution by the sample confidence $\mathbf{c}$ only. It is thus straightforward to incorporate robustness into existing ML implementations.

2.3. Robust SR with postprocessing noise reduction

Since the Gaussian error norm handles inliers just like the quadratic norm does, the performance of the Robust ML (RML) deconvolution is similar to that of the ML deconvolution for normally distributed noise. This is validated in figure 3b and 3c, which show comparable deconvolution results of a Gaussian blurred and noisy image (figure 3a, $\sigma_{psf} = 1$, $\sigma_n = 5$) after five iterations. However, when five percent of salt-and-pepper noise is added to the blurred and noisy image (figure 3d), overshoots start to appear in the ML estimate (figure 3e). The result of the RML deconvolution in figure 3f, on the other hand, remains sharp without outliers. Experiments with more salt-and-pepper noise reveal that the RML deconvolution can handle as much as 25% outliers before it fails to produce a satisfactory result. There is one requirement that must be fulfilled to achieve this performance: the initial estimate $\hat{\mathbf{z}}^0$ must be free of outliers. This is the reason why a 3×3 median filter is applied to the shift&add image in step 1 of the robust SR algorithm in table 1.

Step 1:	Begin at $n = 0$ with an initial estimate $\hat{\mathbf{z}}^0$ being a 3×3 median filtered image of a shift & add solution [4].
Step 2:	Compute $\mathbf{g}^0$ using eq. (13), and initialize the conjugate-gradient vector as $\mathbf{d}^0 = -\mathbf{g}^0$.
Step 3:	Compute the optimal step size ε^n using eq. (17).
Step 4:	Let $\hat{\mathbf{z}}^{n+1} = \hat{\mathbf{z}}^n + \varepsilon^n \mathbf{d}^n$ and $\xi^n = \lVert\hat{\mathbf{z}}^{n+1} - \hat{\mathbf{z}}^n\rVert / \lVert\hat{\mathbf{z}}^n\rVert$.
Step 5:	If $\xi^n < T$ or $\xi^n > \xi^{n-1}$, let $\hat{\mathbf{z}}_{RML} = \hat{\mathbf{z}}^{n+1}$ and stop.
Step 6:	Compute $\mathbf{g}^{n+1}$ and let $\mathbf{d}^{n+1} = -\mathbf{g}^{n+1} + \beta^n \mathbf{d}^n$, where $\beta^n = \Big((\mathbf{g}^{n+1})^T \mathbf{g}^{n+1}\Big) \Big/ \Big((\mathbf{g}^n)^T \mathbf{g}^n\Big)$.
Step 7:	Let $n \leftarrow n + 1$ and go to step 3.

Table 1. A conjugate-gradient optimization algorithm for robust SR restoration.

(a) blurry&noisy→RMSE=10.7 (b) ML of (a) → RMSE=10.9 (c) RML of (a) → RMSE=10.7

(d) (a)+outliers→RMSE=34.1 (e) ML of (d) → RMSE=31.8 (f) RML of (d) → RMSE=11.1

Figure 3. First row: robust ML deconvolution produces a similar result as ML deconvolution for a blurred ($\sigma_{psf} = 1$) and noisy input ($N(0, \sigma_n = 5)$). Second row: the ML result is hampered by outliers (5% salt & pepper noise) whereas the robust ML result is insensitive to them.

A poor response of the ML restoration against outliers is the main reason why it is not preferred over regularized methods such as the Maximum A Posteriori (MAP) restoration [8]. We solved this problem by replacing the quadratic error norm with a robust error norm. Another drawback of the ML solution is the amplification of noise due to which the solution finally converges to a noisy result. Although the iterative process can be stopped before the inverse solution becomes too noisy, there are always some noise present such as the visible graininess in figure 3f. A common solution to penalize these fluctuations is adding a smoothness term $\Upsilon(\hat{\mathbf{z}})$ to the restoration cost function:

$$\hat{\mathbf{z}}_{MAP} = \arg\min_{\mathbf{z}} \left[\|\mathbf{y} - \mathbf{W}\mathbf{z}\|_2^2 + \lambda \Upsilon(\mathbf{z}) \right] \tag{18}$$

Some popular regularization functions are Tikhonov-Miller (TM) [13], Total Variation (TV) [21], and Bilateral Total Variation (BTV) [5]:

$$\begin{aligned} \Upsilon_{TM}(\mathbf{z}) &= \|\Gamma\mathbf{z}\|_2^2 \\ \Upsilon_{TV}(\mathbf{z}) &= \|\nabla\mathbf{z}\|_1 \\ \Upsilon_{BTV}(\mathbf{z}) &= \underbrace{\sum_{l=-2}^{2}\sum_{m=0}^{2}}_{l+m\geq 0} \alpha^{|m|+|l|} \left\| \mathbf{z} - S_x^l S_y^m \mathbf{z} \right\|_1 \end{aligned} \tag{19}$$

where Γ is a Laplacian operator, ∇ is a gradient operator, and S_x^l is a shift operator along x-dimention by l pixels. However, smoothing regularization comes at a cost of signal variance attenuation. The extra term creates a new problem of balancing between the two cost functions (via the regularization parameter λ) [10]. Furthermore, the computational cost of minimizing the combined cost function is usually high.

Instead of minimizing a combined cost function of a data misfit and a data smoothing term, we choose to suppress noise by a postprocessing step. This approach is efficient because the image update per iteration is much simpler without the smoothing term. Although the denoising step seems extraneous, its purpose is equivalent to that of the data smoothing term in the restoration equation (18). Postprocessing denoise is applied to the SR image, which usually contains much fewer pixels than the total number of input LR pixels to be processed by the iterative regularization. The decoupling of image restoration and image denoising allows the flexibility of using many available noise reduction algorithms and easily fine-tuning their parameters. In this section, we use the method of iterated TV refinement [15]. This method was shown to outperform other TV-based and non-TV-based denoising methods for a variety of noise sources. Deconvolution followed by denoising is not new in the literature. Li and Santosa [11] used a similar approach for their computationally efficient TV restoration. ForWaRD [14] (Fourier-Wavelet Regularized Deconvolution) is another deblurring method that separates noise reduction from deconvolution.

The result of deconvolution with iterative TV postprocessing is presented in figure 4 together with the results of other methods. The same salt-and-pepper noisy image from figure 3d is used as input to all tested algorithms. The parameters to all methods are chosen to minimize the Root Mean Square Error (RMSE[1]) of the output and the original image in figure 3c. As can be seen in figure 4a, the Tikhonov-Miller deconvolution [8] with an optimal regularization parameter [10] is not robust against outliers. Its solution is also smooth because of a quadratic regularization norm. Strongly influenced by outliers, the ForWaRD method [14] does not yield a good result either. Although bilateral TV deconvolution is robust, its result in figure 4d is slightly more blurred and noisier than our RML+TV result as depicted in figure 4e. Our RML+TV deblurred image also scores the smallest RMSE against the ground truth. This strong performance is partially due to the excellent noise reduction and signal preservation of

---

[1] $RMSE(f,g) = \sqrt{\frac{1}{N}\sum(f-g)^2}$, where N is the number of samples in f, g

the iterative TV denoising algorithm. As can be seen in figure 4f, the noise residual removed by the iterative TV method shows very little signal leakage.

(a) Tikhonov-Miller→RMSE=24 (b) ForWaRD → RMSE=18.1 (c) original 8-bit image

(d) bilateral TV ($\lambda = 10^{-2}$, optimal step size) → RMSE=9.8 (e) Robust ML + iterative TV denoising → RMSE=8.6 (f) noise removed by TV denoising, linear stretch from -10 to 10

Figure 4. Robust deconvolution followed by iterative TV denoising outperforms other deconvolution methods due to its robustness against outliers. The noise residual removed by TV denoising shows only a minimal signal leakage.

3. Evaluations of super-resolution algorithms

Many SR algorithms have been proposed during the past twenty years, but little has been done to compare their performance. The most commonly used methods for comparison are still the Mean Square Error (MSE) and visual inspection by a human observer. The MSE, applicable only when the ground truth is available, is a questionable measure. A small MSE also does not always correspond to a better image quality [24]. Visual inspection, on the other hand, is dependent not only on the viewers but also on the displays and viewing conditions. In this section, we propose a range of objective measures to evaluate the results of different SR algorithms. These quantitative measures are presented in a radar chart per SR image so that their strengths and weaknesses are easily comparable.

3.1. Objective performance measures for SR

The main goal of SR is to improve the image resolution and to reduce noise and other artifacts. As a result, we propose the following performance measures: the SNR, image sharpness, image jaggedness, size and height (contrast) of a smallest available blob. All these quantities are measured directly from the SR reconstructed image. Together with other criteria such as the absence of artifacts, the performance of SR can be objectively evaluated.

The measurement of these quality factors are illustrated on the SR result in figure 5. Figure 5a shows one of 98 LR input images under translational motion. The SR image in figure 5b is reconstructed from the LR images using the Hardie method [8] with automatic selection of the regularization parameter [10]. Edge sharpness and jaggedness are measured from the sharpest edge in the image (enclosed in a rectangle). A small blob close to the edge is used to measure how a SR algorithm can retain small features. Due to its nonrobustness, the Hardie algorithm fails to remove dead-pixel artifacts from the image. These artifacts are even more visible in figure 5c, where nonlinear histogram equalization is applied to enhance the image contrast.

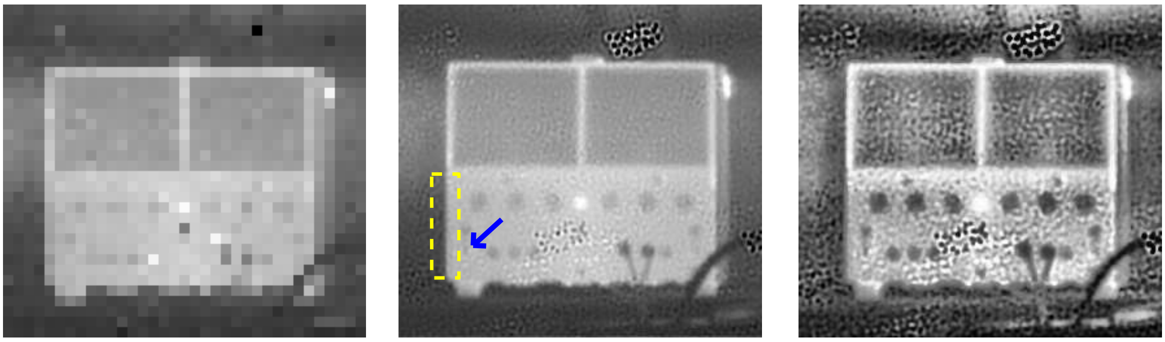

(a) Pixel replication

(b) 5 iterations of Hardie ($\sigma_{psf} = 1.5$, $\lambda = 6.46 \times 10^{-6}$)

(c) nonlinear stretch of (b)

Figure 5. Eight-times zoom from 98 LR images. The rectangle and the arrow point to the edge and the blob used in later analysis. (a) and (b) are linearly stretched between [18400 20400], (c) is stretched by adaptive histogram equalization [26].

3.1.1. Signal-to-noise ratio To measure the SNR[2] of a SR image, the energies of the signal and the noise are required. Since the ground truth HR image is not available, we estimate the noise variance from a Gaussian fit to the histogram of a masked high-pass filtered image. The high-pass filtered image in figure 6a ($\sigma_{hp} = 3$) is segmented into discontinuities (lines, edges, and texture) and homogeneous regions by a threshold on its absolute intensities. The segmentation in figure 6b is computed from the least noisy SR result (by the Zomet method), and the same mask is used for all other SR images. After the noise variance is estimated, the variance of the noise-free signal is estimated as a difference of the *corrupted* signal variance and the noise variance.

3.1.2. Edge sharpness and jaggedness Another desirable quality of a super-resolved image is sharpness. We measure the sharpness as the width of the Edge Spread Function (σ_{esf}) across a straight edge [20]. Together with the average edge width, the location of the straight edge is also determined with subpixel accuracy. Edge jaggedness, on the other hand, is measured as the standard deviation of the estimated edge positions along the straight edge. This position jaggedness is most likely due to registration errors between the LR frames. We determine the edge crossing point of each 1-D profile perpendicular to the edge by fitting an $error\ function$

---

[2] $SNR = 10\log_{10}(\sigma_I^2/\sigma_n^2)$, where σ_I^2 and σ_n^2 are energies of the noise-free signal and noise, respectively.

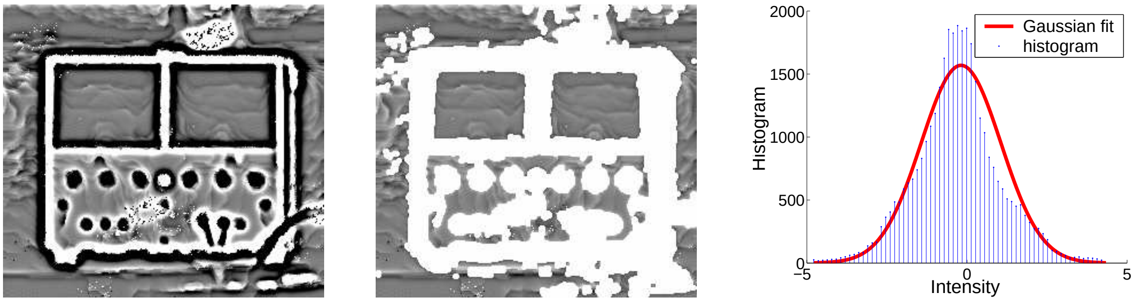

(a) high-pass filtered image (b) non-textured mask (c) Gaussian fit to noise histogram

Figure 6. Noise variance estimation from a single image by a Gaussian fit to the histogram of homogeneous regions in a high-pass filtered image ($\sigma_{hp} = 3 \longrightarrow \sigma_n = 1.25$)

(erf) to the profile (see figure 7d):

$$\hat{x}_0 = \arg\min_{x_0} \left\| I(x) - A.erf\left(\frac{x - x_0}{\sigma_{esf}}\right) - B \right\|_1 \tag{20}$$

where x_0 is the edge crossing point, x is the coordinate along the 1-D profile, $I(x)$ is the profile intensity, $erf(z) = \frac{2}{\pi}\int_0^z e^{-t^2}dt$ is the error function, A and B are intensity scaling and offset, respectively. The standard deviation of these estimated edge crossing points $\hat{x}_0$ from the fitted straight edge is defined as the edge jaggedness.

Figure 7 illustrates a step-by-step procedure to compute edge sharpness and jaggedness. Figure 7a shows an elevation view of the Region Of Interest (ROI) enclosed by the dashed box in figure 5b. Because of shading, the intensity levels on either side of the edge are not flat. This shading is corrected line-by-line, which results in a normalized step edge in figure 7b. 2-D Gaussian blur edge fit and individual 1-D erf fit are applied to the normalized edge. The results are depicted in figure 7c, in which the straight edge is superimposed as a thick dashed line, and the horizontal edge crossing points are connected by a thin dark line. Because the jaggedness of the dark line is negligible compared to the edge width (stdev(x_0) $= 0.12 << \sigma_{esf} = 2.79$), this edge is quite smooth. Figure 7d also shows that the edge profile can be modeled very accurately by an error function.

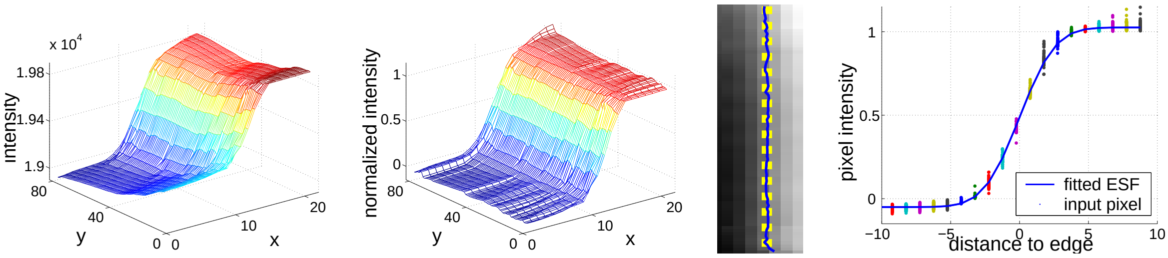

(a) edge from Zomet result (b) normalized edge (c) line fit (d) Gaussian edge profile

Figure 7. Finding the sharpness and jaggedness of a straight edge ($\sigma_{esf} = 2.79 \pm 0.12$). The edge is zoomed four-times horizontally in (c), in which the thick dashed line is the fitted edge and the thin jagged line connects all rows' inflection points.

3.1.3. Reconstruction of small details In addition to edges, small details are also of interest in many SR tasks. We measure the size and contrast of the smallest detectable blob in the SR image by fitting a 2-D Gaussian profile to its neighborhood. Figure 8a shows part of the apparatus that is pointed to by the arrow in figure 5a. This ROI contains two blobs of respectively 3.2 and 2.0 HR pixels radius. The bigger blob is visible in the results of all tested SR algorithms (figures 5 and 9). However, the contrast of the smaller blob is in some cases too weak to be seen. The 3-D mesh views of the two blobs and their fitted Gaussian surfaces are plotted in figure 8b and 8c on the same scale. The Gaussian profile fits well to the image intensities in both cases. The amplitude (or height) of the smaller blob is about half that of the bigger blob. This can be verified in figure 8a, in which the contrast of the smaller blob is significantly weaker than that of the bigger one.

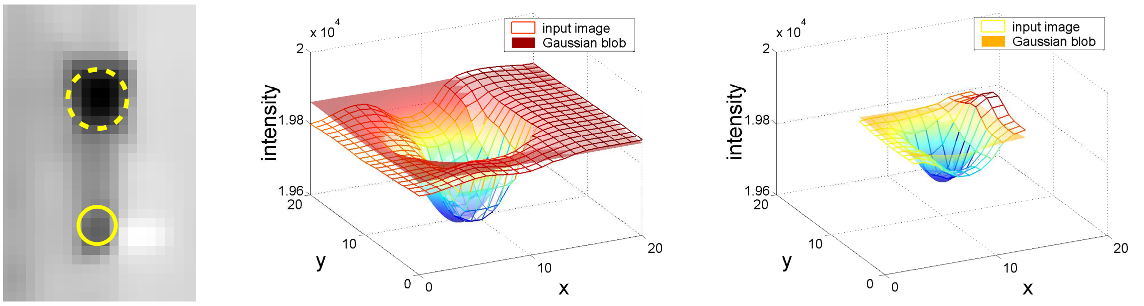

(a) SR result (RNC+ML+BF)

(b) Top blob: $\sigma = 3.2$, height=277

(c) Bottom blob: $\sigma = 2.0$, height=132

Figure 8. Fitting of 2-D Gaussian profiles to small blobs in a SR image.

3.2. Evaluation results

In this subsection, the results of four robust SR methods are compared against each others based on the metrics of the presented performance measures. The first two methods from [5] are the regularized Zomet method, which iteratively updates the SR image by a median error back projection, and the Farsiu method, which is essentially a Bilateral Total Variation (BTV) deconvolution of a shift-and-add fusion image. The remaining two methods are from the authors of this paper. The RNC+ML+BF method follows a separate reconstruction-restoration approach as the Farsiu method. First-order Robust Normalized Convolution (RNC) [18] is used for the reconstruction, Maximum-Likelihood (ML) deconvolution (section 2.1) is used for the restoration, and an xy-separable Bilateral Filtering (BF) [17] postprocessing step is used for noise reduction. The RML+BF method, on the other hand, follows a direct multi-frame restoration approach as the Zomet method. Instead of using the median operator to achieve robustness, the Robust Maximum Likelihood (RML) restoration (section 2.2) uses a local mode estimation to reduce the outliers' influence in the back projection. The RML+BF method is also followed by a bilateral filtering step for noise removal.

Our robust SR methods differ from most other methods in the literature in the separate treatment of normally distributed noise and outliers. We do not use a data regularization term for outlier removal because heavy regularization might be required, which in turn produces an overly smoothed result. Instead, the outliers are suppressed by penalizing their influences during the RNC fusion or the RML restoration. In addition, while the Zomet and the Farsiu methods require iterative regularization for noise suppression, we only need one single pass of bilateral filtering after restoration. This postprocessing strategy is more flexible because the filter parameters or even the denoising algorithms can be changed without rerunning the lengthy restoration process. Optimal selection of the regularization parameter during the iterative restoration is also difficult. Too much regularization in the first several iterations biases the solution to a smooth result. Too little regularization in the last several iterations produces a noisy output image. With a single pass of bilateral filtering, the parameters only need to be initialized once (e.g., $\sigma_s = 3$, $\sigma_t = 3\sigma_n$). The only concern to take into account is that the iterative ML restoration should be stopped before noise is amplified beyond the contrast level of interested details.

3.2.1. The apparatus sequence The results of the four SR algorithms on the apparatus sequence are presented in figure 9. The Zomet and Farsiu results are produced using the software from UCSC[3]. A Gaussian PSF with $\sigma_{psf} = 1.5$ HR pixel is used for all methods. Other parameters of the Zomet and Farsiu methods including the regularization parameter λ and the step size β are manually fine-tuned to produce a visually best result in terms of noise reduction and details preservation. To keep the small blob on the left side of the apparatus visible, outliers are not removed completely from the Zomet and the Farsiu results. The parameters of our SR methods, on the other hand, are automatically determined from the noise level in the image. The tonal scales of RNC fusion, RML restoration, and BF denoising, for example, are chosen approximately three-times the standard deviation of noise before the operations. Due to a separate treatment of outliers and Gaussian noise, both of our SR results remain artifact-free while true details are well preserved.

The performance measures are presented in a radar chart format for each SR result in figure 9. Each radar axis corresponds to one of the five performance measures: standard deviation of noise, edge sharpness, edge jaggedness, blob size and an inverse of blob height (the selected edge and blob are shown earlier in figure 5b). Note that the blob height is inverted in the last performance measure so that a smaller value corresponds to a better quality in all axes. Although the origins of all axes start from the common center, each axis has a different unit and scale. The axes are scaled by their average measures over the four given images. These average measures are shown as a dotted equilateral pentagon in all charts. A set of five performance measures per image forms a pentagon which is shown in thick lines. Although it is desirable to have a thick pentagon for which all measures are the smallest, such a clear winner is not present in figure 9. A more realistic way to evaluate different SR algorithms is to compare their relative pentagon areas. Because parameter tuning can only reduce a certain measures at the cost of increasing others (e.g., noise versus blob size), the area of the pentagon does not vary significantly for a specific algorithm. As a result, an algorithm with the smallest pentagon area generally yields the best performance.

Based on the SR images and the performance radar charts in figure 9, the strengths and weaknesses of each SR method are revealed. The Zomet+BTV regularization method, for example, produces a low-noise output but it fails to preserve small details. Outlier pixels are still present because the pixel-wise median operator is only robust when there are more than three LR samples per HR pixels (the ratio is only 1.5 in this case). The Farsiu method is not robust to this kind of extreme outliers either. Those outlier spots in figure 9c spread out even more compared to those in the Zomet result due to the L_2 data norm within the Farsiu L_2+BTV deconvolution. The BTV regularization is, however, good at producing smooth edges. Both the Zomet and the Farsiu methods score a below-average edge jaggedness measure.

While the regularized methods have difficulties in balancing between outlier removal and detail preservation, our methods based on sample certainty produce outlier-free outputs with a high level of detail preservation. Both the RNC+ML+BF and the RML+BF methods reconstruct small blobs very well. However, the standard deviations of noise are below average.

---

[3] MDSP resolution enhancement software: `http://www.ee.ucsc.edu/~milanfar`

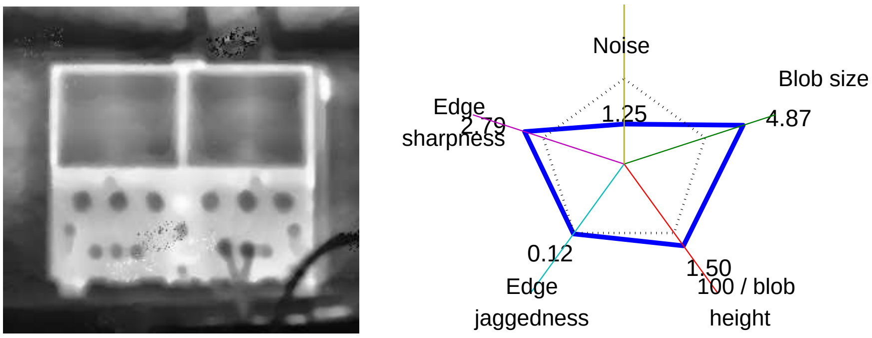

(a) Zomet+BTV ($\lambda = 0.002$, $\beta = 5$) (b) SNR = 42.5 dB, blob height = 67

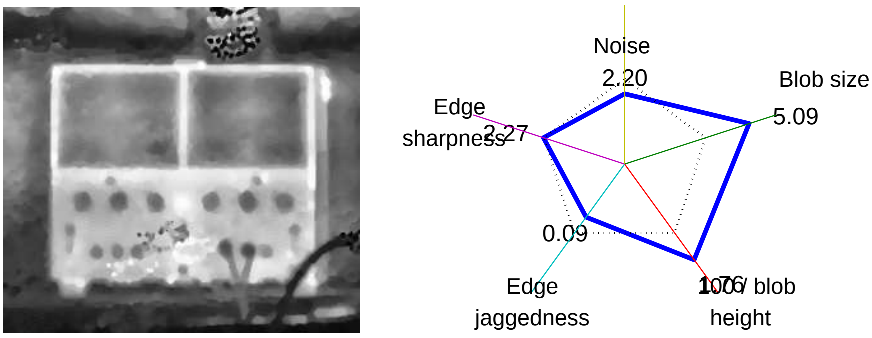

(c) Farsiu ($\lambda = 0.002$, $\beta = 20$) (d) SNR = 37.8 dB, blob height = 57

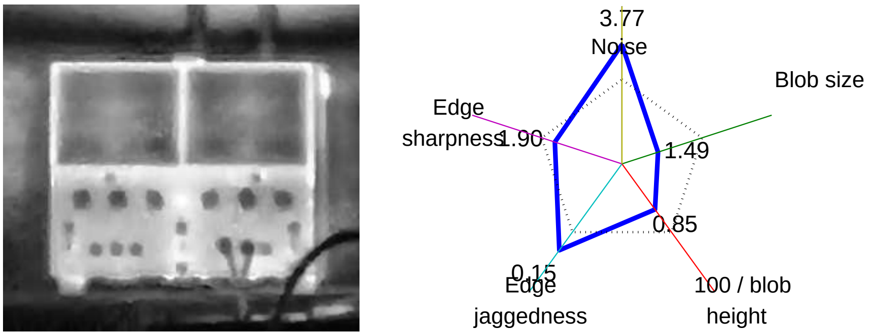

(e) RNC ($\sigma_s = 1$, $\sigma_t = 50$) + ML + BF ($\sigma_s = 3$, $\sigma_t = 30$) (f) SNR = 40.7 dB, blob height = 119

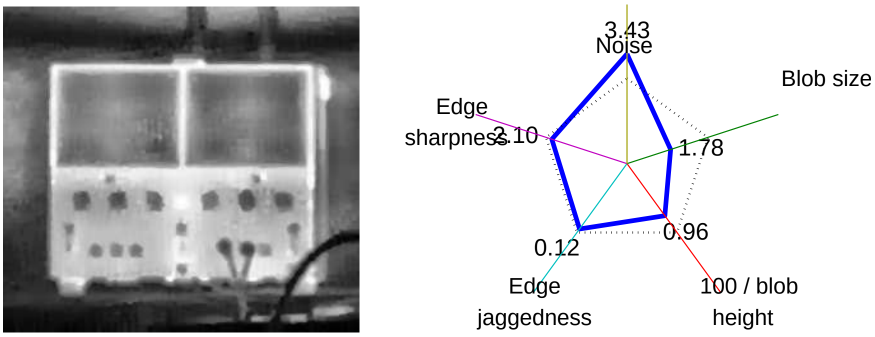

(g) RML (σ_t=55) + BF (σ_s=3, σ_t=30) (h) SNR = 41.5 dB, blob height = 117

Figure 9. Eight-times SR from 98 shifted LR images. All images are nonlinearly stretched [26]. The average performance measures (dotted) are [2.66 3.31 1.27 0.12 2.26]. Lower values indicate better performance.

Considering that the SNR is large in all images ($SNR \approx 40$), this is not a real problem. The RML+BF method seems to produce the best result with four out of five measures being better than average. Visually speaking, its result is also better than the RNC+ML+BF result because the latter shows some slight ringing artifacts near the top and the bottom of the apparatus.

3.2.2. The bar chart sequence To show that the presented performance measures are not tied to a single example, the same evaluation process is applied to a different image sequence. The infrared sequence in figure 10a contains 98 LR images captured by a manually-induced jitter-moving camera. Apart from random noise and aliasing, the LR images are also corrupted by outliers, which are caused by permanently dead and blinking sensor elements. Using the robust RML+BF method, a clean SR result is reconstructed in figure 10b. Four-times upsampling is used in this experiment because a higher zoom does not produce more details.

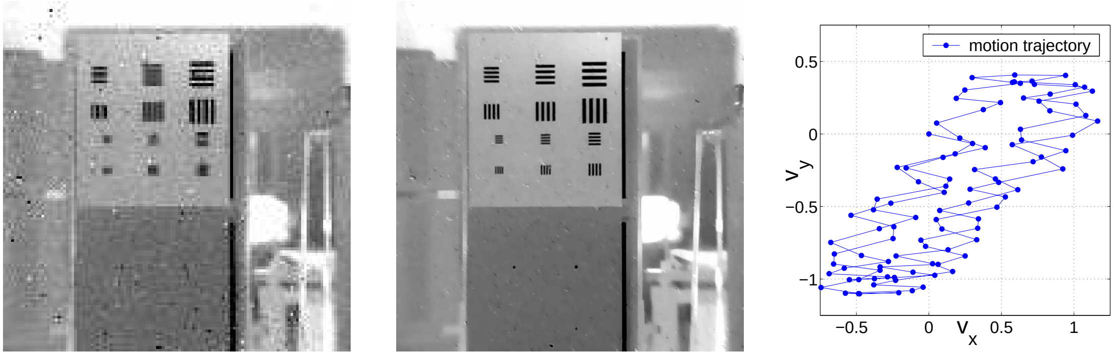

(a) 1 of 98 LR image (128×128) (b) 4× SR by RML+BF (c) Motion trajectory

Figure 10. Four-times SR of a shifted sequence by the robust ML restoration followed by bilateral filtering. The images are linearly stretched between [18200 19100].

A zoomed-in version of the resolution chart can be seen in figure 11. Due to low resolution and aliasing, the parallel bars in figure 11a are not separable. They are, however, clearly visible in figure 11b after a Hardie SR reconstruction. Although the Hardie SR result reveals more details than the LR image, it also shows spotty artifacts due to its non-robust nature. The directed artifact is the trace of a single outlier pixel along the sequence's motion trajectory (figure 10c). Similar to the previous experiment, a step edge at the bottom of the plate and a small blob at the top-right of the plate are used for performance measurements. The mask in figure 11c shows non-textured regions that are used for noise estimation in all subsequent SR results.

The results and analysis of four SR algorithms: Zomet, Farsiu, RNC+ML+BF, and RML+BF, applied to the resolution chart sequence are presented in figure 12. Similar to the previous example, the parameters for the Zomet and the Farsiu methods are manually selected so that all parallel bars are clearly distinguishable. The parameters for the RNC+ML+BF and the RML+BF methods, on the other hand, are chosen automatically as described in [18] and section 2.1, respectively. As can be seen from the left column of figure 12, outlier spots are present in all four SR images. Each spot is formed by the same outlier pixel under the subpixel motion of the sequence ($|v_x|$, $|v_y|$ are less than one LR pixel). These outliers are not completely removed by the median operation in the Zomet method. The Farsiu method, which performs a deconvolution on the non-robust Shift & Add (SA) fusion image, is the most noisy of the four results because these large outlier spots are not treated as outliers by the L_2+BTV deconvolution algorithm. Our last two results show the best noise reduction of the four methods

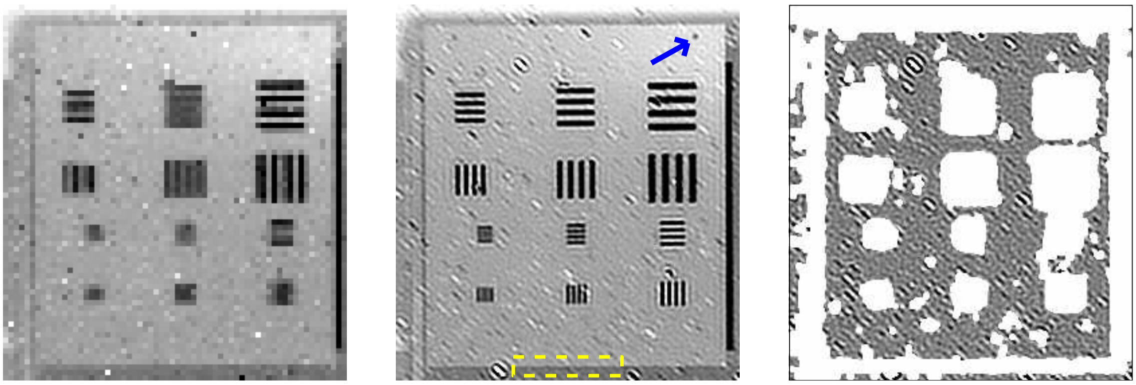

(a) Pixel replication (b) Hardie (σ_{psf}=1, λ=4×10^{-6}) (c) noise mask → $\sigma_n = 30.2$

Figure 11. Four-times SR using a non-robust method together with the edge, blob, and noise mask used for later quality assessment. (a) and (b) are linearly stretched between [18200 19100], (c) is stretched between [-128 128].

thanks to the sample certainty approach. Apart from a lower noise level, our SR methods also reconstruct sharper edges and finer blobs. The RNC+ML+BF method produces the least noise and the lowest edge jaggedness. The RML+BF method scores better than average on all five performance measures. These quantitative results are strongly supported by visual inspection of the accompanying images.

3.2.3. The eye chart sequence The last sequence to be evaluated in this subsection is shown in figure 13. Aliasing is visible in the lower part of the LR image in figure 13a as some vertical lines are not as clear as others. Four-times SR by a robust method produces a HR image without outliers in figure 13b. Because this sequence moves with a larger shifts than that of the bar chart sequence, outliers spread out more evenly in the HR grid and they are therefore removed more thoroughly (see figure 15 for the results of four robust SR methods).

Noise, aliasing and intensity outliers are clearly visible in a zoomed-in version of the LR image in figure 14a. After a four-times SR by the Hardie methods, the openings of the circles are resolvable. Figure 14b also shows the step edge and the small blob used in subsequent performance analysis. Figure 14c shows a noise mask, which is composed of only flat areas for noise estimation.

The performance of different SR algorithms applied to the eye chart sequence is presented in figure 15. The Farsiu method performs worst in terms of noise and detail contrast. Although a better SNR is possible by more regularization, this solution only further smears away the smallest circles. Due to a sufficient number of LR samples (six LR samples per HR pixel) and a good separability of the outlier samples after registration, the Zomet method performs reasonably well on this sequence. All its performance measures are around the average. However, due to the data smoothing term, both the Zomet and the Farsiu results are not as sharp as ours in the last two figures. Visually speaking, the RML+BF result is the sharpest, although this advantage is slightly hampered by an above average edge jaggedness. This last sequence eventually shows that given many LR samples per HR pixel and a negligible amount of outliers, the difference in performance of the four presented methods are subtle. Non-blind evaluation methods, such as the triangle orientation discrimination method [2], may be necessary to resolve these differences.

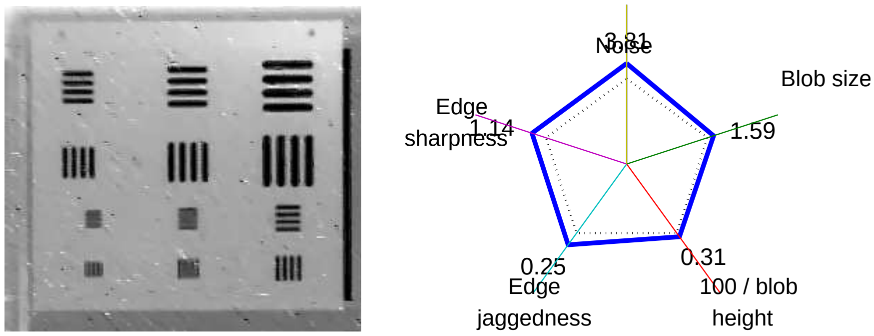

(a) Zomet+BTV (λ=0.005, β=10) (b) SNR = 32.4 dB, blob height = 323

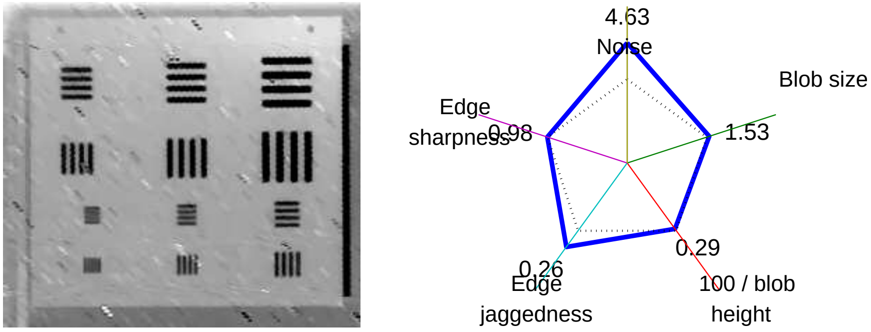

(c) Farsiu ($\lambda = 0.01$, $\beta = 10$) (d) SNR = 30.9 dB, blob height = 350

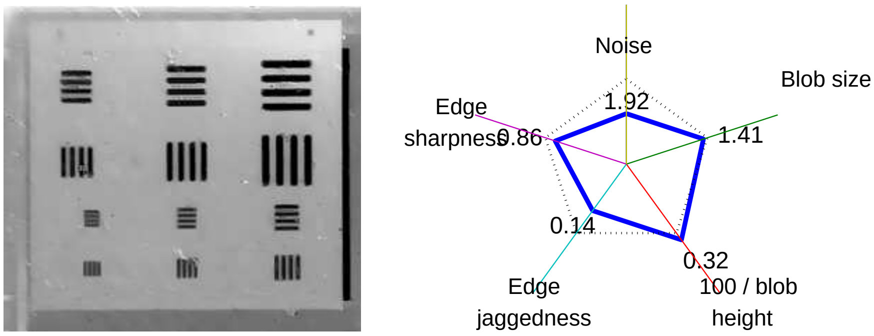

(e) RNC (σ_s=0.5, σ_t=50) + ML (6 iters) + BF ($\sigma_s = 3$, $\sigma_t = 50$) (f) SNR = 38.7 dB, blob height = 310

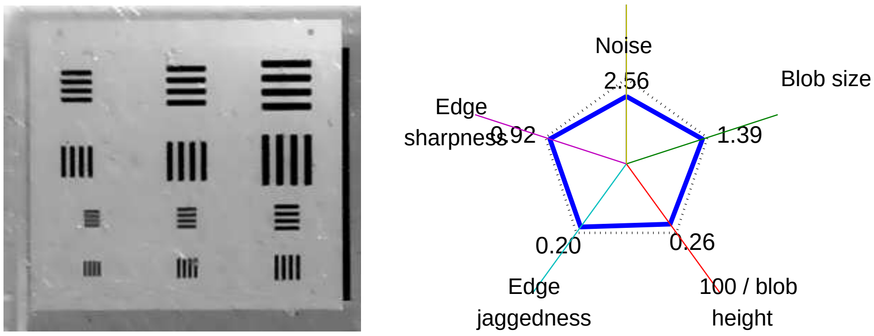

(g) RML ($\sigma_t = 55$, 20 iters) + BF ($\sigma_s = 3$, $\sigma_t = 60$) (h) SNR = 36.6 dB, blob height = 390

Figure 12. Four-times SR from 98 shifted LR images using different algorithms. The average performance measures are [3.23 1.48 0.29 0.21 0.98] (the dotted pentagon). Lower values indicate a higher performance.

(a) 1 of 98 LR image (128×128) (b) 4× SR by RML+BF (c) Motion trajectory

Figure 13. Four-times SR of a shifted sequence by the robust ML restoration followed by bilateral filtering. The images are linearly stretched between [18000 19000].

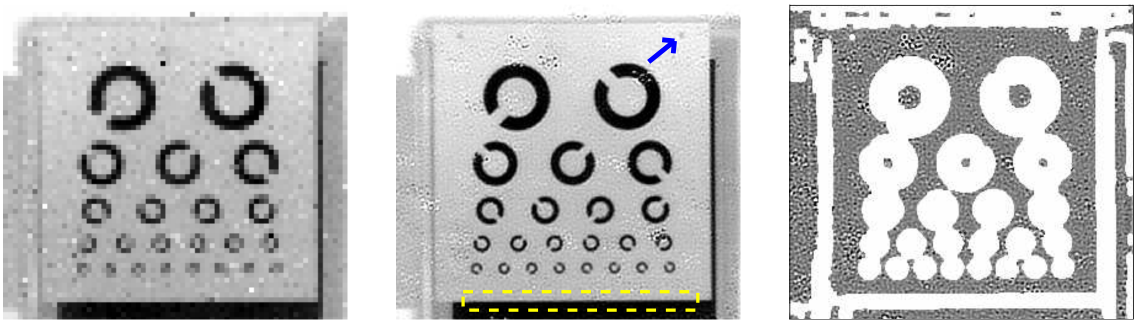

(a) Pixel replication (b) Hardie (σ_{psf}=1,λ=3.7×10^{-6}) (c) noise mask → σ_n = 29.1

Figure 14. Four-times SR using the Hardie method together with the edge (dashed rectangle), blob (pointed to by an arrow), and noise mask used for quality assessment.

4. Conclusion

We have presented a new super-resolution restoration algorithm that is robust against outliers. Our robust algorithm differs from the standard ML algorithm by a tonal-based sample certainty component. Noise reduction is implemented as a postprocessing step instead of a regularization term in the cost function of the iterative restoration. Due to a simpler cost function, our SR algorithm has a lower computational complexity than the regularized methods. Noise is also reduced more effectively without compromising the signal quality. Evaluations of four robust SR algorithms on three different datasets have confirmed the superior performance of our algorithm over other methods in the literature.

Another contribution of this paper is a scheme for objective comparison of different SR outputs. Provided that edges and blobs are present, the evaluation scheme first computes five quality factors: standard deviation of the noise, edge sharpness, edge jaggedness, blob size, and blob contrast. It then scales and plots these measurements on a radar chart, producing a polygon representation of the measures for each image. Because a smaller value means a better quality, the polygon with a smaller area generally corresponds to a better SR output. Different from other evaluation schemes in the literature, our scheme is objective and it produces a consistent result over a number of tested datasets.

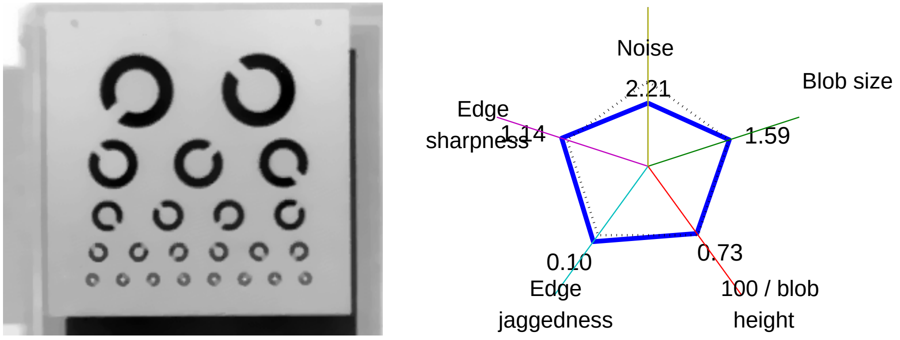

(a) Zomet+BTV (λ=0.02, β=5) (b) SNR = 40.6 dB, blob height = 137

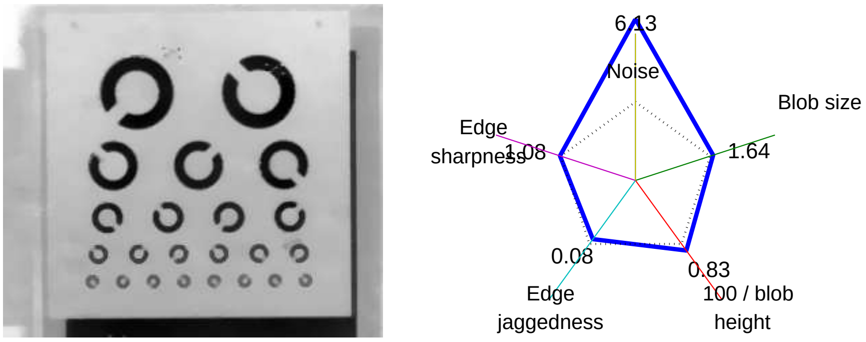

(c) Farsiu ($\lambda = 0.05$, $\beta = 5$) (d) SNR = 31.7 dB, blob height = 121

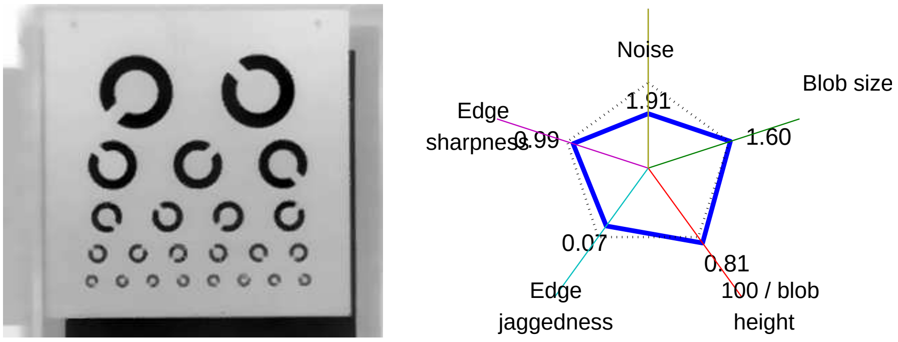

(e) RNC (σ_s=0.5, σ_t=60) + ML (6 iters) + BF ($\sigma_s = 3$, $\sigma_t = 50$) (f) SNR = 42.3 dB, blob height = 123

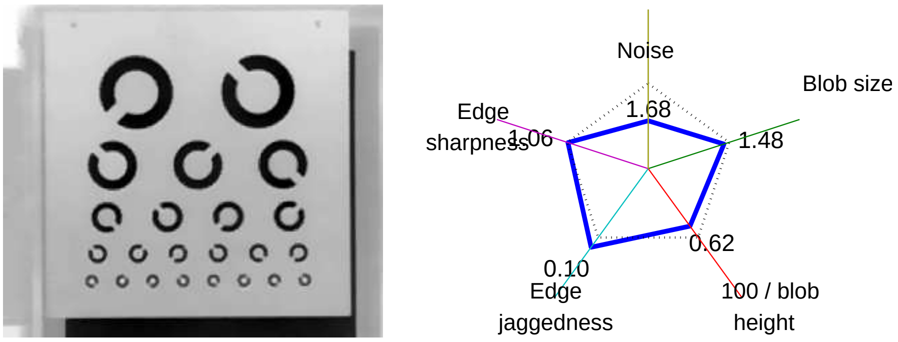

(g) RML ($\sigma_t = 57$, 10 iters) + BF ($\sigma_s = 3$, $\sigma_t = 60$) (h) SNR = 43.5 dB, blob height = 160

Figure 15. Four-times SR from 98 shifted LR images using different algorithms. The average performance measures are [2.98 1.58 0.75 0.09 1.07] (the dotted pentagon). Lower values indicate a higher performance.

References

[1] R. van den Boomgaard and J. van de Weijer. Least squares and robust estimation of local image structure. *IJCV*, 64(2/3):143–155, September 2005.

[2] A.W.M. van Eekeren, K. Schutte, O.R. Oudegeest, and L.J. van Vliet. Performance evaluation of super-resolution reconstruction methods on real-world data. *EURASIP Journal on Applied Signal Processing*, 2007:Article ID 43953, 2006.

[3] M. Elad and A. Feuer. Restoration of a single superresolution image from several blurred, noisy, and undersampled measured images. *IEEE Trans. on IP*, 6(12):1646–1658, 1997.

[4] S. Farsiu, D. Robinson, M. Elad, and P. Milanfar. Robust shift and add approach to super-resolution. In A.G. Tescher, *Applications of Digital Image Processing XXVI*, volume 5203 of *SPIE*, pages 121–130, 2003.

[5] S. Farsiu, M.D. Robinson, M. Elad, and P. Milanfar. Fast and robust multiframe super resolution. *IEEE Trans. on IP*, 13(10):1327–1344, October 2004.

[6] R. Fletcher and C.M. Reeves. Function minimization by conjugate gradients. *The Computer Journal*, 7(2):149–154, 1964.

[7] F.R. Hampel, E.M. Ronchetti, P.J. Rousseeuw, and W.A. Stahel. *Robust Statistics: The Approach Based on Influence Functions*. Wiley, New York, 1986.

[8] R.C. Hardie, K.J. Barnard, J.G. Bognar, E.E. Armstrong, and E.A. Watson. High resolution image reconstruction from a sequence of rotated and translated frames and its application to an infrared imaging system. *Opt. Eng.*, 37(1):247–260, January 1998.

[9] M. Irani and S. Peleg. Improving resolution by image registration. *GMIP*, 53:231–239, 1991.

[10] M.G. Kang and A.K. Katsaggelos. General choice of the regularization functional in regularized image-restoration. *IEEE Trans. on IP*, 4(5):594–602, 1995.

[11] Y.Y. Li and F. Santosa. A computational algorithm for minimizing total variation in image restoration. *IEEE Trans. on IP*, 5(6):987–995, 1996.

[12] D.G. Luenberger. *Linear and Nonlinear Programming*. Addison-Wesley, 1984.

[13] K. Miller. Least-squares method for ill-posed problems with a prescribed bound. *SIAM J. Math. Anal.*, 1(10):52–74, February 1970.

[14] R. Neelamani, H. Choi, and R. Baraniuk. ForWaRD: Fourier-wavelet regularized deconvolution for ill-conditioned systems. *IEEE Trans. Signal Processing*, 52(2):418–433, 2004.

[15] S. Osher, M. Burger, D. Goldfarb, J. Xu, and W. Yin. An iterative regularization method for total variation-based image restoration. *Multiscale Model. Simul.*, 4(2):460–489, 2005.

[16] S.C. Park, M.K. Park, and M.G. Kang. Super-resolution image reconstruction: A technical overview. *IEEE Signal Processing Magazine*, 20(3):21–36, 2003.

[17] T.Q. Pham and L.J. van Vliet. Separable bilateral filtering for fast video preprocessing. In *CD-ROM proc. of ICME'05*, 2005.

[18] T.Q. Pham, L.J. van Vliet, and K. Schutte. Robust fusion of irregularly sampled data using adaptive normalized convolution. *EURASIP Journal on Applied Signal Processing*, 2006:Article ID 83268, 2006.

[19] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, Cambridge, 2nd edition, 1992.

[20] S. Reichenbach, S. Park, and R. Narayanswamy. Characterizing digital image acquisition devices. *Opt. Eng.*, 30(2):170–177, 1991.

[21] L. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Phys. D*, 60:259–268, 1992.

[22] A.N. Tikhonov and V.Y. Arsenin. *Solutions of Ill-posed problems*. Wiley, Washington, 1977.

[23] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *Proc. of ICCV'98*, pages 839–846, 1998.

[24] Z. Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Trans. on IP*, 13(4):600–612, April 2004.

[25] A. Zomet, A. Rav-Acha, and S. Peleg. Robust super-resolution. In *Proc. of CVPR'01*, pages I:645–650, Kauai, Hawaii, 2001.

[26] K. Zuiderveld. Contrast limited adaptive histogram equalization. In P.S. Heckbert, *Graphics Gems IV*, pages 474–485. Boston Academic Press, 1994.