Robust Fusion of Irregularly Sampled Data Using Adaptive Normalized Convolution

Tuan Q. Pham,[1] Lucas J. van Vliet,[1] and Klamer Schutte[2]

[1] *Quantitative Imaging Group, Department of Imaging Science and Technology, Faculty of Applied Sciences,*
  Delft University of Technology, Lorentzweg 1, 2628 CJ, Delft, the Netherlands
[2] *Electro Optics Group, TNO Defence, Security, and Safety, P.O. Box 96864, 2509 JG, the Hague, the Netherlands*

We present a novel algorithm for image fusion from irregularly sampled data. The method is based on the framework of normalized convolution (NC), in which the local signal is approximated through a projection onto a subspace. The use of polynomial basis functions in this paper makes NC equivalent to a local Taylor series expansion. Unlike the traditional framework, however, the window function of adaptive NC is adapted to local linear structures. This leads to more samples of the same modality being gathered for the analysis, which in turn improves signal-to-noise ratio and reduces diffusion across discontinuities. A robust signal certainty is also adapted to the sample intensities to minimize the influence of outliers. Excellent fusion capability of adaptive NC is demonstrated through an application of super-resolution image reconstruction.

1. INTRODUCTION

In digital image processing, continuous signals are often digitized on a regular grid. Data in this form greatly simplifies both hardware design and software analysis. As a result, if an image is available in another format, it is often resampled onto a regular grid before further processing. Super-resolution (SR) reconstruction of shifted images under common space-invariant blur, in particular, reconstructs a high-resolution (HR) image from a set of randomly positioned low-resolution (LR) images. While there are many approaches that achieve SR through an iterative minimization of a criterion function [12, 13, 30], this paper is concerned with SR fusion as a separate step after image registration and before deblurring.

A popular method for fusion of irregularly sampled data is surface interpolation. A triangulation-based method [15], for example, first computes a Delaunay tessellation of the data points, then interpolates the data locally within each tile. The triangulation method, aiming to be an exact surface interpolator, is not designed to handle noisy data. It is also expensive to tessellate in achieving SR because of the large number of LR samples involved. Though computationally less expensive, other surface interpolation methods, such as the inverse distance-weighted method and the radial basis function method [1], are all very sensitive to noise.

In the presence of noise, a surface fit is often preferred over exact interpolation. A polynomial approximation to a small neighborhood in the image, known as the facet model, has been proposed by Haralick as early as 1981 [11]. The Haralick facet model, however, is not well localized for large neighborhoods since all data points have equal importance. Farnebäck [7] corrects this by introducing a Gaussian applicability to the operator, which puts more emphasis on fitting the central pixels. van den Boomgaard and van de Weijer [27] further extend the facet model with a robust error norm to handle a mixture of models around image discontinuities. However, none of these facet models are explicitly designed for irregularly sampled data, which requires a sample localization mechanism like the Delaunay triangulation [15]. Another drawback of these methods is that they ignore the fact that natural images are often comprised of directional structures, and that the image derivatives can be integrated along these structures to improve their estimation.

In this paper, we introduce a robust certainty and a structure-adaptive applicability function to the polynomial facet model and apply it to fusion of irregularly sampled data. The method is based on normalized convolution (NC) [14], in which the local signal is approximated through a projection onto a subspace spanned by a set of basis functions. Unlike the traditional framework, however, the operator's applicability function adapts to local linear structures. This leads to more samples of the same modality being gathered for the analysis, which in turn improves signal-to-noise ratio (SNR) and reduces diffusion across discontinuities. The

robust signal certainty is incorporated to minimize the influence of outliers caused by dead pixels or occasional misregistration.

The paper is organized as follows. Section 2 reviews the idea of normalized convolution and its least-squares solution. Section 3 introduces robustness to NC via a robust signal certainty. The certainty is estimated directly from the intensity difference between the current sample and its neighbors. Section 4 presents a rotated anisotropic Gaussian applicability function. The steering parameters for the adaptive applicability function are computed from gradient information of the input data. An example on real infrared images in Section 5 shows that excellent SR reconstruction with high SNR is achievable with image fusion using the robust and adaptive NC.

2. NORMALIZED CONVOLUTION USING POLYNOMIAL BASES

Normalized convolution (NC) [14] is a technique for local signal modeling from projections onto a set of basis functions. Although any bases can be used, the most common one is a polynomial basis: $\{\mathbf{1}, \mathbf{x}, \mathbf{y}, \mathbf{x}^2, \mathbf{y}^2, \mathbf{xy}, \ldots\}$, where $\mathbf{1} = [1 \ 1 \ \cdots \ 1]^T$ (N entries), $\mathbf{x} = [x_1 \ x_2 \ \cdots \ x_N]^T$, $\mathbf{x}^2 = [x_1^2 \ x_2^2 \ \cdots \ x_N^2]^T$, and so on are constructed from local coordinates of N input samples. The use of polynomial basis functions make the traditional NC equivalent to a local Taylor series expansion. Within a local neighborhood centered at $\mathbf{s}_0 = \{x_0, y_0\}$, the intensity value at position $\mathbf{s} = \{x + x_0, y + y_0\}$ is approximated by a polynomial expansion:

$$
\begin{aligned}
\widehat{f}(\mathbf{s}, \mathbf{s}_0) = \ & p_0(\mathbf{s}_0) + p_1(\mathbf{s}_0)x + p_2(\mathbf{s}_0)y + p_3(\mathbf{s}_0)x^2 \\
& + p_4(\mathbf{s}_0)xy + p_5(\mathbf{s}_0)y^2 + \cdots,
\end{aligned}
\tag{1}
$$

where $\{x, y\}$ are the local coordinates of sample $\mathbf{s}$ with respect to the center of analysis $\mathbf{s_0}$. $\mathbf{p}(\mathbf{s}_0) = [p_0 \ p_1 \cdots p_m]^T(\mathbf{s}_0)$ are the projection coefficients onto the corresponding polynomial basis functions at $\mathbf{s}_0$.

Different from the Haralick facet model [11], which is also a polynomial expansion, NC uses a so-called applicability function to localize the polynomial fit (while the facet model gives an equal weight to all samples in a neighborhood). This applicability function is often an isotropic, radially decaying function whose size is proportioned to the scale of analysis. A Gaussian function is often used for this purpose. The projection $\mathbf{p}(\mathbf{s}_0)$ can then be used to derive Gaussian derivatives, which are image projections onto Hermite polynomials [28]. In addition, NC allows each input signal to have its own certainty value. The signal certainty is especially useful when data samples are missing or are unreliable (e.g., due to bad sensors or erroneous registration). Both the applicability function and the signal certainty control the impact of a particular sample to the local polynomial fit.

The choice of the polynomial order depends on specific applications. If processing speed is more important than accuracy, NC with a constant basis is sufficient. This locally flat model, however, does not model edges and ridges very

well. First-order NC with three bases $\{\mathbf{1}, \mathbf{x}, \mathbf{y}\}$ can model edges, and second-order NC with six bases $\{\mathbf{1}, \mathbf{x}, \mathbf{y}, \mathbf{x}^2, \mathbf{xy}, \mathbf{y}^2\}$ can further model ridges and blobs. Higher-order NC can fit more complex structures at a higher computational cost. However, NC with order greater than two is rarely used since the high-order bases are often fit to noise rather than the signal itself. In this paper, we propose to use first-order NC for SR fusion.

The scale of the applicability function also plays a decisive role in the quality of interpolation. Low-order NC with a large applicability window cannot reconstruct small details in the image. The scale of the applicability function, however, must be large enough to cover sufficient samples for a stable local analysis. Unless the sample density is high everywhere in the image (e.g., in case of SR from many LR frames), a normal choice of the applicability function is a Gaussian function with a spatial scale $\sigma_s = 1$ HR pixel and a truncation of three standard deviations. This Gaussian applicability function introduces minimal blurring to the interpolation result while its support is still large to cover enough samples.

2.1. Least-squares estimation

To solve for the projection coefficients $\mathbf{p}$ at an output position $\mathbf{s}_0$, the approximation error is minimized over the extent of an applicability function a centered at $\mathbf{s}_0$:

$$
\varepsilon(\mathbf{s}_0) = \int \left(f(\mathbf{s}) - \widehat{f}(\mathbf{s}, \mathbf{s}_0) \right)^2 c(\mathbf{s}) \, a(\mathbf{s} - \mathbf{s}_0) \, d\mathbf{s}, \tag{2}
$$

where the signal certainty $0 \le c(\mathbf{s}) \le 1$ specifies the reliability of the measurement at $\mathbf{s}$, with zero representing completely untrustworthy data and one representing very reliable data. Although both c and a act as scalar weights for the squared errors, they represent different properties, each of which can be made adaptive to the local image data as shown in the next two sections. For a neighborhood encompassing N samples, standard least-squares regression yields a solution in matrix form [7]:

$$
\mathbf{p} = \left(\mathbf{B}^T \mathbf{W} \mathbf{B} \right)^{-1} \mathbf{B}^T \mathbf{W} \mathbf{f}, \tag{3}
$$

where $\mathbf{f}$ is an $N \times 1$ matrix of input intensity $f(\mathbf{s})$, $\mathbf{B} = [\mathbf{b}_1 \ \mathbf{b}_2 \cdots \mathbf{b}_m]$ is an $N \times m$ matrix of m basis functions sampled at local coordinates of N input samples, and $\mathbf{W} = \mathrm{diag}(\mathbf{c})$. $\mathrm{diag}(\mathbf{a})$ is an $N \times N$ diagonal matrix constructed from an element-by-element product of the signal certainty $\mathbf{c}$ and the *sampled* applicability $\mathbf{a}$.

In case of regularly sampled data with a fixed certainty and a fixed applicability function, NC can be implemented very efficiently by convolution operations only. Since the local neighborhood is organized in the same regular grid, the basis functions are also fixed. The least-squares solution in (3) for zeroth-order NC can be simplified to two convolutions:

$$
\widehat{f}_0 = \frac{a \otimes (c \cdot f)}{a \otimes c}, \tag{4}
$$

where $\widehat{f}_0$ is the interpolated image, $\otimes$ is the convolution operator, and $c \cdot f$ is the pixel-wise multiplication of the certainty
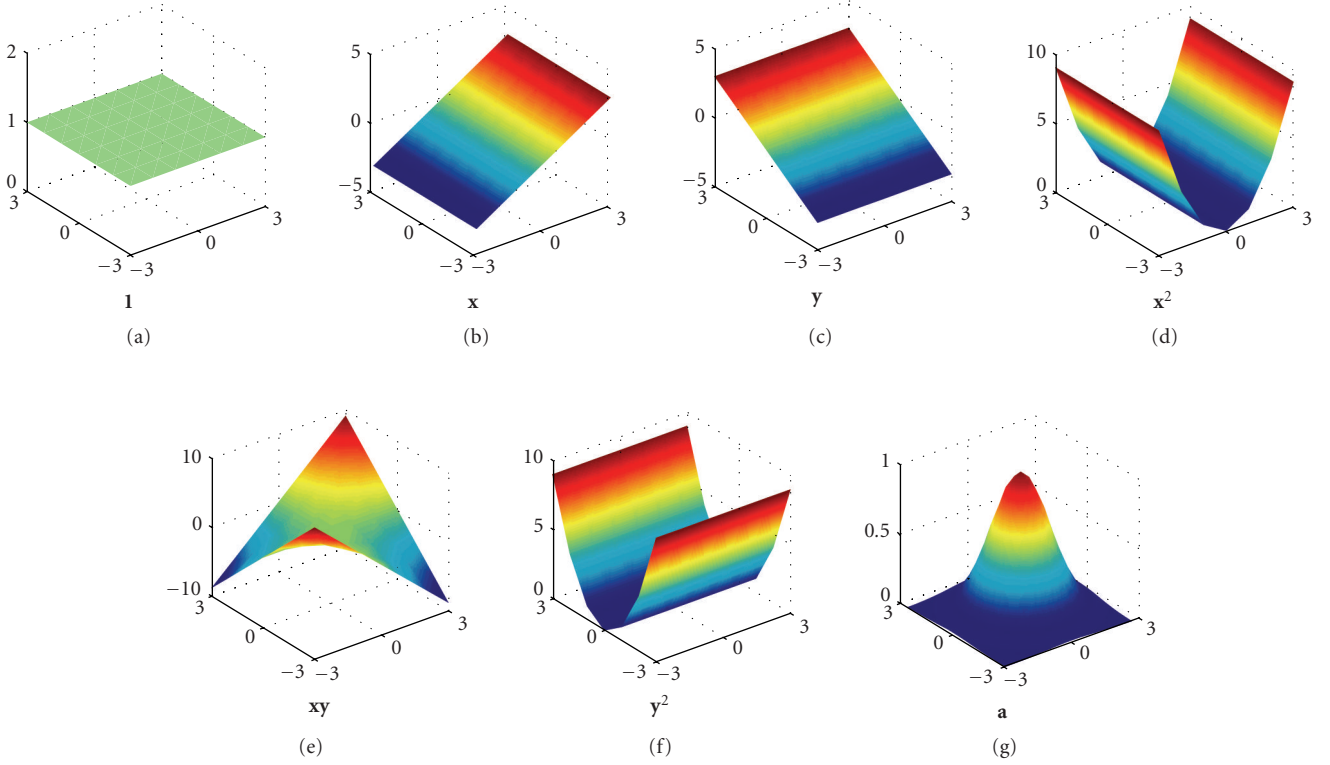
FIGURE 1: Polynomial basis functions $\{\mathbf{1}, \mathbf{x}, \mathbf{y}, \mathbf{x}^2, \mathbf{xy}, \mathbf{y}^2\}$ and Gaussian applicability function $\mathbf{a}$.

image and the intensity image. A full first-order NC requires nine convolutions and produces three output images: an interpolated image $\widehat{f}_1$ and two directional derivatives $\widehat{f}_x$, $\widehat{f}_y$ in the x- and y-dimensions:

$$
\begin{bmatrix} \widehat{f}_1 \\ \widehat{f}_x \\ \widehat{f}_y \end{bmatrix} = \left(\begin{bmatrix} a & a.x & a.y \\ a.x & a.x^2 & a.xy \\ a.y & a.xy & a.y^2 \end{bmatrix} \otimes c \right)^{-1}
$$
$$
\times \left(\begin{bmatrix} a \\ a.x \\ a.y \end{bmatrix} \otimes (c \cdot f) \right),
$$
(5)

where x, y, x^2, xy, y^2, and a are two-dimensional kernels of the basis functions and applicability function as shown in Figure 1. NC on a regular grid can be spedup even further by separable and recursive convolution [29] if a Gaussian applicability function is used. The denominator in (4) and the matrix inversion in (5) are normalization terms to correct for the nonhomogeneous signal certainty, hence the name normalized convolution.

2.2. Irregular sample collection

Unfortunately, NC does not reduce to a set of regular convolutions for irregularly sampled signals because the polynomial bases and applicability functions are sampled at irregular local coordinates. Each output position therefore requires

a different matrix multiplication and inversion. Moreover, since the samples are irregularly positioned, they must first be gathered before a local analysis.

To ensure a fast local sample collection, we setup a reference list at each pixel on a regular output grid to keep records of input samples within half a pixel away. These data structures are initialized once before fusion. They can shrink or grow as samples are removed or added. This is useful for dynamic super-resolution of video where new frames are inserted and old frames are removed from the system. To gather all samples within several pixels away from an output position, the references are collected from the records stored at all grid points in the neighborhood. Since it is easier to traverse through a regular grid than a set of irregular points, input samples can be collected more efficiently with these reference lists. The data structure, though simple, provides a tremendous saving of sample searching time. It is also compact because only the references are kept rather than all sample attributes.

Irregular sample collection could be done more efficiently in the case of SR fusion of shifted LR frames with an integer zoom factor. If the zoom factor μ is an integer, the pattern of LR sample distribution is repetitive after each $\mu \times \mu$ pixel block in the HR grid. Provided that the applicability function is fixed, the reference lists should only be constructed for μ^2 pixels in the first $\mu \times \mu$ image block. Every other output pixel at coordinates $\{x, y\}$ then takes the same local sample organization as the pixel at $\{x - \mu \lfloor x/\mu \rfloor, y - \mu \lfloor y/\mu \rfloor\}$
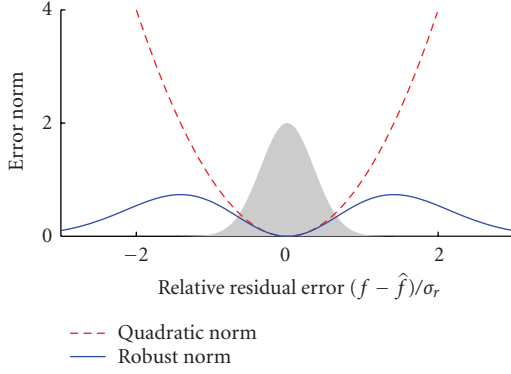
FIGURE 2: Robust norm $\Psi(f, \hat{f}) = |f - \hat{f}|^2 \exp(-|f - \hat{f}|^2/2\sigma_r^2)$ versus quadratic norm $\Psi(f, \hat{f}) = |f - \hat{f}|^2$.

in the first block (where $\lfloor \cdot \rfloor$ is the integer floor operator and $x - \mu\lfloor x/\mu \rfloor$ is the remainder of the division of x by μ). The same local sample organization here means the local samples come from the same LR frames but at a $\{\lfloor x/\mu \rfloor, \lfloor y/\mu \rfloor\}$ offset in LR pixels. In this way, the applicability $a(\mathbf{s} - \mathbf{s}_0)$ could be precomputed for all irregular sample $\mathbf{s}$ around $\mathbf{s}_0$, leading to an efficient implementation of (3).

3. ROBUST NORMALIZED CONVOLUTION

While NC is a good interpolator for uncertain data, it requires the signal certainty to be known in advance. With the same photometric-based weighting scheme used in bilateral filtering [24], a robust certainty is assigned to each neighboring sample before a local polynomial expansion around $\mathbf{s}_0$. The robust certainty, being a Gaussian function of residual error $f - \hat{f}$, assigns low weights to potential outliers, effectively excluding them from the analysis:

$$c(\mathbf{s}, \mathbf{s}_0) = \exp\left(-\frac{|f(\mathbf{s}) - \hat{f}(\mathbf{s}, \mathbf{s}_0)|^2}{2\sigma_r^2}\right), \qquad (6)$$

where $f(\mathbf{s})$ is a measured intensity at position $\mathbf{s}$ and $\hat{f}(\mathbf{s}, \mathbf{s}_0)$ is an estimated intensity at $\mathbf{s}$ using an initial polynomial expansion at the center of analysis $\mathbf{s}_0$. Unlike the fixed certainty $c(\mathbf{s})$ in (2) that depends only on the position $\mathbf{s}$, the robust certainty $c(\mathbf{s}, \mathbf{s}_0)$ changes as the window of analysis moves. The photometric spread σ_r defines an acceptable range of the residual error $f - \hat{f}$. Samples with residual error less than σ_r get a certainty close to one, whereas those with residual error larger than $2 \times \sigma_r$ get an extremely low certainty. We select σ_r to be two times the standard deviation of input noise (σ_{noise} is estimated from low-gradient regions in the image) so that all samples within $\pm 2\sigma_{\text{noise}}$ deviation from the initial polynomial surface fit get a certainty close to one.

The product of a quadratic norm $|f - \hat{f}|^2$ and the Gaussian certainty in (6) results in an error norm that is robust against outliers. Figure 2 compares this robust norm with a quadratic norm. While the quadratic norm keeps increasing at higher residual error, the robust norm peaks at a residual

error of $\sqrt{2}\sigma_r$; it then reduces to practically zero for large residual error. The shaded profile in this figure shows a typical Gaussian distribution of the inlier residual. Since the photometric spread σ_r is chosen to be twice larger than the noise spread σ_{noise}, the robust norm behaves like a quadratic norm for all normally distributed noise; it then gradually reduces to zero outside $\pm 3\sigma_{\text{noise}}$ to reject outliers. With this adaptive certainty, NC becomes a weighted least-squares estimator that behaves as a normal least-squares estimator under Gaussian noise and it is robust against outliers.

One problem remains with robust NC: it does not have a closed-form solution as in the case of least-squares NC. Due to the certainty (6), the robust polynomial expansion requires an initial estimation of the polynomial expansion itself. However, similar to the analysis of bilateral filtering in [5, 27], robust NC can be solved by an iterative weighted least-squares minimization. Started with an initial polynomial expansion (we use a flat model at a locally weighted median [3] level), the certainty can be computed according to (6). The weighted least-squares estimation is then solved by (3), resulting in an updated polynomial expansion. The process is repeated until convergence (three iterations are often enough). It has been shown in [25] that this iterative procedure quickly converges to a closest local maximum of a local histogram observed at a spatial scale σ_s and a tonal scale σ_r, a.k.a. the local mode. Initialization that is close to the true intensity is therefore crucial. Although the weighted median is generally a robust choice as an initial estimate, the closest sample is sometimes used instead. The latter is applicable in image filtering when noise level is low or when minute details are of interest after filtering.

The impact of the robust certainty on NC fusion of data with outliers can be seen in Figure 3. In this experiment, ten LR images are generated from the HR image in Figure 3(a) by randomly shifting the original image followed by three-time downsampling in both directions. The LR images are then corrupted by five percent of salt and pepper noise, one of them is shown in Figure 3(b). Four fusion methods[1] are applied to the data: L2 regularized back-projection by Hardie [12], L2 data norm with bilateral total-variation regularization (L2 + bilateral TV) by Farsiu [9], robust fusion using median of back-projected errors by Zomet [30], and our robust NC. The parameters for these methods are tuned for a smallest root mean-squared Error between the reconstructed and the original image:

$$\text{RMSE}(f, \hat{f}) = \sqrt{\frac{1}{N} \sum (f - \hat{f})^2}, \qquad (7)$$

where N is the number of samples in $f, \hat{f}$. Fifty iterations are used for the three methods [9, 12, 30] because it takes that many iterations for the methods to converge with this highly contaminated data. Since the Hardie method is not designed

---

[1] Implementations of [9, 30] are available with a Matlab toolbox at http://www.ee.ucsc.edu/~milanfar.
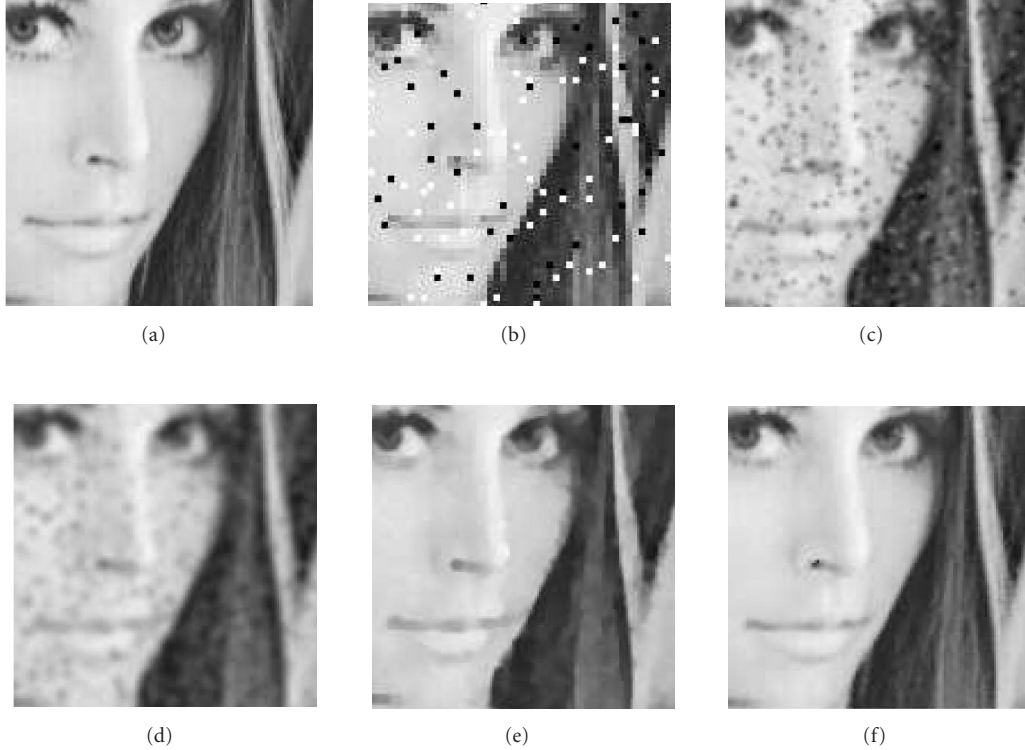
FIGURE 3: Three-times upsampling of 10 shifted LR images corrupted by 5% salt and pepper noise. The parameter settings were obtained by minimizing the RMSE. (a) Original 8-bit image; (b) 1 of 10 LR inputs + 5% salt and pepper noise → RMSE = 12.3; (c) Hardie conjugate gradient [12], $\lambda = 8.3$ → RMSE = 14.6; (d) Zomet [30] + L2 regularize $\lambda = 0.15$, $\beta = 5$ → RMSE = 10.2; (e) Farsiu L2 + bilateral TV [9] $\lambda = 0.15$, $\beta = 1.68$, $\sigma_{PSF} = 1.24$ → RMSE = 7.4; and (f) robust first-order NC, $\sigma_s = 0.6$, $\sigma_r = 10$ → RMSE = 6.5.

for robustness, a large regularization parameter ($\lambda = 8.3$) is required to suppress the salt and pepper noise. Yet, too much regularization smoothens the image while noise is not completely removed (Figure 3(c)). The iterative robust fusion methods do not perform well on this high level of outliers either. While the Zomet method produces good reconstruction for less than one percent outliers,[2] it breaks at five percent salt and pepper noise. The blurred output in Figure 3(d) is a fusion result of Zomet method with norm 2 regularization parameter $\lambda = 0.15$ and a step size $\beta = 5$. The Farsiu method ($\lambda = 0.16$, $\beta = 1.78$, and a Gaussian deconvolution kernel at scale $\sigma_{PSF} = 1.24$) successfully removes all outliers but the result looks cartoon-like due to the TV regularization. Furthermore, because the same regularization used to remove outliers is applied to uncorrupted pixels, small details are not reconstructed very well by L2 + bilateral TV. Our result using robust NC removes most of the outliers after only two iterations compared to 50 iterations of other methods. Small details such as irises, eyelashes, and hair pieces are well reconstructed by robust NC. An analysis of the RMSE between the reconstructed and the original image also confirms superior performance of robust NC over the other methods.

---

[2] Experiments were done but the results are not shown here.

4. STRUCTURE-ADAPTIVE NORMALIZED CONVOLUTION

NC is a local operator in a sense that it requires a finite neighborhood to operate. First-order NC in 2D, for example, requires at least three samples to fit a local plane. If there are ample samples per pixel, the scale of the applicability function could be very small, leading to a sharp image reconstruction. However, in underdetermined cases where input samples are sparse, the applicability scale must be increased to gather enough samples for a stable polynomial fit at the expense of a blurrier result. However, an applicability function that only extends along linear structures will not diffuse across lines and edges. Therefore, the edge-enhanced fusion result stays sharp for the purpose of small detail perception. In this section, we present such an adaptive applicability function and show that it significantly increases the quality of sparsely sampled data interpolation.

We use a spatially adaptive filtering kernel similar to that of Nitzberg and Shiota [17]. The applicability function is an anisotropic Gaussian kernel that adapts its shape and orientation along the underlying image structure. The adaptive applicability function ensures that only samples sharing similar intensity and gradient information are gathered for the local polynomial expansion. The kernel is extended along the
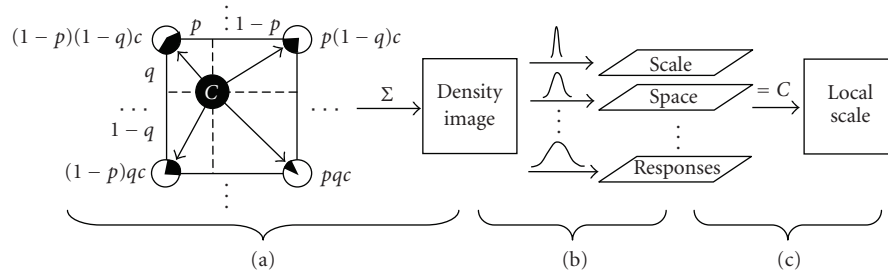
FIGURE 4: Fast estimation of local scale by a quadratic interpolation along the scale axis of a Gaussian scale-space of the HR density image.

local linear structure allowing better noise suppression while avoiding signal blurring across lines and edges. Since samples along a linear structure share similar gradient information, the adaptive applicability function is applicable to an NC of any order.

4.1. Estimation of local image structure and scale

To construct an adaptive kernel at an output pixel, the local image structure around that pixel must be known in advance. We compute an initial estimate of the output intensity I and gradient information $I_x = \partial I / \partial x$ and $I_y = \partial I / \partial y$ using first-order robust NC from the previous section. Local structure information including orientation ϕ and anisotropy A is computed from the eigenvectors $\{\mathbf{u}, \mathbf{v}\}$ and the corresponding eigenvalues ($\lambda_u \geq \lambda_v$) of a principal component analysis of the local gradient vectors $\nabla I = [I_x \ I_y]^T$ (a.k.a. the gradient structure tensor (GST) method) [26]:

$$\text{GST} = \overline{\nabla I \, \nabla I^T} = \begin{bmatrix} \overline{I_x^2} & \overline{I_x I_y} \\ \overline{I_x I_y} & \overline{I_y^2} \end{bmatrix} = \lambda_u \mathbf{u}\mathbf{u}^T + \lambda_v \mathbf{v}\mathbf{v}^T,$$
$$\phi = \arg(\mathbf{u}), \qquad A = \frac{\lambda_u - \lambda_v}{\lambda_u + \lambda_v}, \tag{8}$$

where the tensor elements are averaged locally by a Gaussian filter at a scale of 1.5 pixels. The tensor smoothing integrates the structural information over several neighboring pixels and is thus less susceptible to noise than the information from a single gradient vector. However, this tensor smoothing also means that the estimated structural information is valid for that particular scale only. As a result, if small features are of interest, a small tensor scale should be used.

Another important data characteristic is local sample density, since it reveals how much information is available near the HR grid points. In the case of uncertain data, the sample density is computed as a sum of sample certainty over an unnormalized Gaussian-weighted neighborhood of scale $\sigma_c(\mathbf{s}_0)$ (i.e., a Gaussian kernel whose middle weight equals one):

$$d(\mathbf{s}_0, \sigma_c) = \sum \exp\left(-\frac{(\mathbf{s}-\mathbf{s}_0)^2}{2\sigma_c^2(\mathbf{s}_0)}\right) c(\mathbf{s}, \mathbf{s}_0). \tag{9}$$

We define a local scale $\sigma_c(\mathbf{s}_0)$ as the scale at which $d(\mathbf{s}_0, \sigma_c)$ is equal to a constant C ($C = 1$ for zero-order NC, $C = 3$ for first-order NC). The size of the applicability function is

then set to this scale to minimize smoothing in regions with high sample density. To estimate this local scale, we use a quick algorithm as depicted in Figure 4. The certainty of each irregular sample is split to its four nearest HR grid points in a bilinear-weighting fashion (Figure 4(a)). The accumulation of all grid-stamped sample certainties forms *a density image* on the HR grid (Figure 4(b)). A Gaussian scale-space of this density image at exponentially increasing scales ($\sigma_i = 2^i$, $i = -1, 0, 1, 2, \ldots$) is constructed using fast separable and recursive filtering [29] (note that the filter weights are not normalized, that is, the maximum filter tap is one). Due to the unnormalized filter weights, the scale-space responses at each pixel increase with a quadratic rate. We can then perform a quadratic interpolation at each grid point along the scale axis to estimate the Gaussian scale whose filter response is equal to C (Figure 4(c)).

4.2. Structure-adaptive applicability function

The adaptive applicability function is an anisotropic Gaussianfunction whose main axis is rotated to align with the local dominantorientation:

$$a(\mathbf{s}, \mathbf{s}_0)$$
$$= \rho(\mathbf{s}-\mathbf{s}_0)\exp\left[-\left(\frac{x\cos\phi+y\sin\phi}{\sigma_u(\mathbf{s}_0)}\right)^2-\left(\frac{-x\sin\phi+y\cos\phi}{\sigma_v(\mathbf{s}_0)}\right)^2\right], \tag{10}$$

where $\mathbf{s}_0 = \{x_0, y_0\}$ is the center of analysis, $\mathbf{s} - \mathbf{s}_0 = \{x, y\}$ are the local coordinates of input samples with respect to $\mathbf{s}_0$. ρ is a pillbox function centered at the origin that limits the kernel support to a certain radius. σ_u and σ_v are the directional scales of the anisotropic Gaussian kernel. σ_v is the scale along the elongated orientation and is greater than or equal to σ_u (see Figure 5). The two directional scales are adjusted by the local scale σ_c estimated in the previous subsection. The local scale σ_c allows the applicability function to shrink or grow depending on how densely populated the neighborhood is:

$$\sigma_u = \frac{\alpha}{\alpha+A}\sigma_c, \qquad \sigma_v = \frac{\alpha+A}{\alpha}\sigma_c. \tag{11}$$

The tuning parameter $\alpha > 0$ sets an upper-bound on the eccentricity of the applicability function (we use $\alpha = 1/2$ for a maximum eccentricity of 3 when the anisotropy $A = 1$). Note that we do not shape the directional filter scale according to the inverse of the eigenvalues of the GST as in [17] to
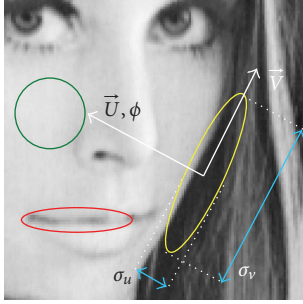
FIGURE 5: Examples of structure-adaptive applicability functions (the scales are exaggerated).

prevent a degeneration of the kernel into an infinitely long ellipse.

Although the computational complexities of all flavors of NC are linear with respect to the number of input samples, robust NC with an isotropic applicability function runs much faster than adaptive NC. This is partly due to the coordinate transformation that takes place under the adaptive scheme. Our implementation of robust NC with the isotropic applicability function is currently two times faster than Matlab's implementation of Delaunay interpolation (griddata.m). With an adaptive applicability function, however, NC of all samples is somewhat slower. Fortunately, since adaptive NC is performed as a second pass after a robust NC, it can be selectively applied to highly anisotropic pixels (pixels with anisotropy $A > 0.5$), whose results could improve significantly from the first pass. This selected fusion saves a lot of computation time without compromising the quality of output signals.

An example of SR fusion for a severely underdetermined case using structure adaptive NC is illustrated in Figure 6. Five input images are generated from the same HR image in the first experiment by randomly shifting the HR image before downsampling five-times in both directions. The generated LR images are then fused together to form a five-times upsampled image. Since there are only five LR images for a zooming factor of five in both directions, the setting is severely underdetermined. Adaptive NC is compared against three iterative methods: Farsiu [9], Zomet [30], and Hardie [12]. The parameter settings for the latter three methods are manually tuned for the smallest RMSE. Visual inspection showed that all of them have converged after about 50 iterations. Even though the original HR image is not blurred before downsampling, both Zomet and Farsiu methods require a deconvolution kernel to produce a sharper image. This is because these algorithms slightly blur its HR image reconstruction when rounding the offsets of input frames to its nearest integer positions on the HR grid. Deconvolution kernel is not used for the Hardie method because it only enhances the jitter artifacts and increases the RMSE. In fact, all iterative methods produce jaggy edges for this underdetermined example because the isotropic regularization does not handle the lack of input samples well. Adaptive NC, on the other hand, reduces the edge jaggedness by extending

its filtering support along linear structures. The images in Figure 6 show that adaptive NC outperforms other fusion methods in terms of both visual quality as well as RMSE. Our method successfully reconstructs the continuation of hair, fur, and hat structures, while other methods simply produce blurred and jittered responses instead.

5. SUPER-RESOLUTION FUSION OF LOW-RESOLUTION IMAGE SEQUENCES

Super-resolution (SR) fusion from a sequence of low-resolution (LR) images is an important step in computer vision toincrease spatial resolution of captured images for subsequent detection, classification, and identification tasks. Extensive literature on this topic exists [2, 4, 6, 9, 12, 13, 15, 23, 30], of which there are two main approaches: one with an integrated fusion and deblurring process [12, 13, 30] and the other with three separate steps: registration, fusion, and deconvolution [6, 9, 15]. The second approach is mostly used when the LR images undergo translational motion and are corrupted by a common space-invariant blur [9].

In this paper, we follow a three-step SR approach as depicted in Figure 7. The LR images are registered against a common frame to a subpixel accuracy using an iterative gradient-based shift estimator [18]. Robust fusion using adaptive NC is then applied to the motion-corrected LR samples. Deconvolution [9] finally reduces the blur and noise caused by optics and sensor integration. The fusion block in Figure 7 is further divided into three substeps, each improving the HR estimate. The first estimate HR_0 is constructed by a locally weighted median operation [3]. HR_0 is then used as an initial estimate for a first-order robust NC, which produces a better estimate of the HR image HR_1 and two derivatives HR_x and HR_y in x- and y-directions. The derivatives are then used to construct anisotropic applicability functions for a final adaptive NC. Implementation details of each fusion substep can be found in the previous sections.

5.1. Super-fusion experiment

In this subsection, a SR experiment is carried out on real data to demonstrate the robust fusion capability of adaptive NC. The input consists of one hundred 128×128 images of a lab scene captured by a pan and tilt camera at long infrared wavelengths (IR with wavelength around $10\,\mu$m). Due to a large pixel pitch with respect to the optical point-spread function (PSF) and a small fill-factor ($\approx 50\%$), the LR images in Figure 8(a) are severely aliased. A resolution enhancement of two in both directions (two-times SR) is therefore possible by fusion alone [20]. With bilateral total variation deconvolution [9], we show that smaller details are resolvable at eight-times SR.

The result of four-time upsampling using adaptive NC for the whole scene is shown in Figure 8(b). The HR image is constructed in the same process as shown in Figure 7. The scale of the applicability function used in the robust NC are $\sigma_u = \sigma_v = 1$ and the photometric spread $\sigma_r = 500$ (1% of the full dynamic range of the 16-bit input images). Two

(a)

(b)

(c)

(d)

FIGURE 6: Five-time edge-enhancing image upsampling from only 20% samples using adaptive NC. (a) Zomet [30] + L1 regularization, $\lambda = 0.001$, $\beta = 2$, $\sigma_{PSF} = 0.8 \rightarrow$ RMSE = 8.2; (b) Farsiu L2 + bilateral TV [9], $\lambda = 0.03$, $\beta = 2$, $\sigma_{PSF} = 0.8 \rightarrow$ RMSE = 7.5; (c) Hardie [12], $\lambda = 1.275 \times 10^{-4} \rightarrow$ RMSE = 7.6; and (d) adaptive zero-order NC $\rightarrow$ RMSE = 6.7.
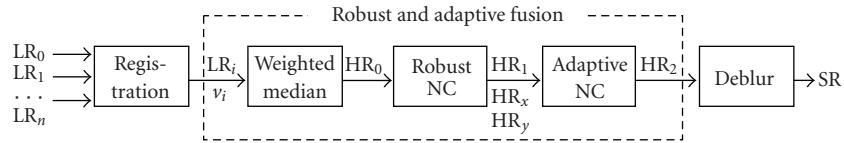

FIGURE 7: Robust and adaptive normalized convolution super-resolution process.
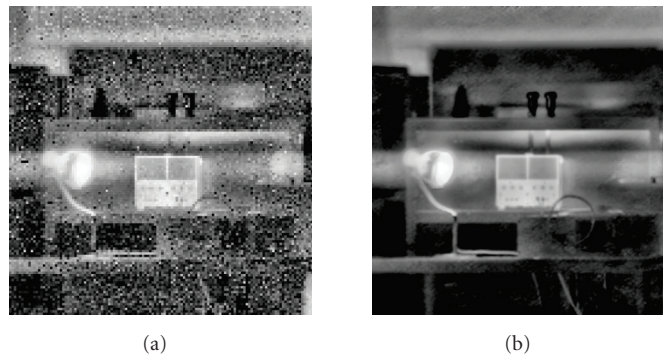

(a)

(b)

FIGURE 8: Four-time increase in resolution of a translated IR sequence by adaptive NC. (The 16-bit images are displayed in 8 bits following an adaptive histogram equalization [31]). (a) 128×128 image captured by a $10\,\mu$m IR camera and (b) $4 \times$ SR fusion from 100 frames by adaptive NC.
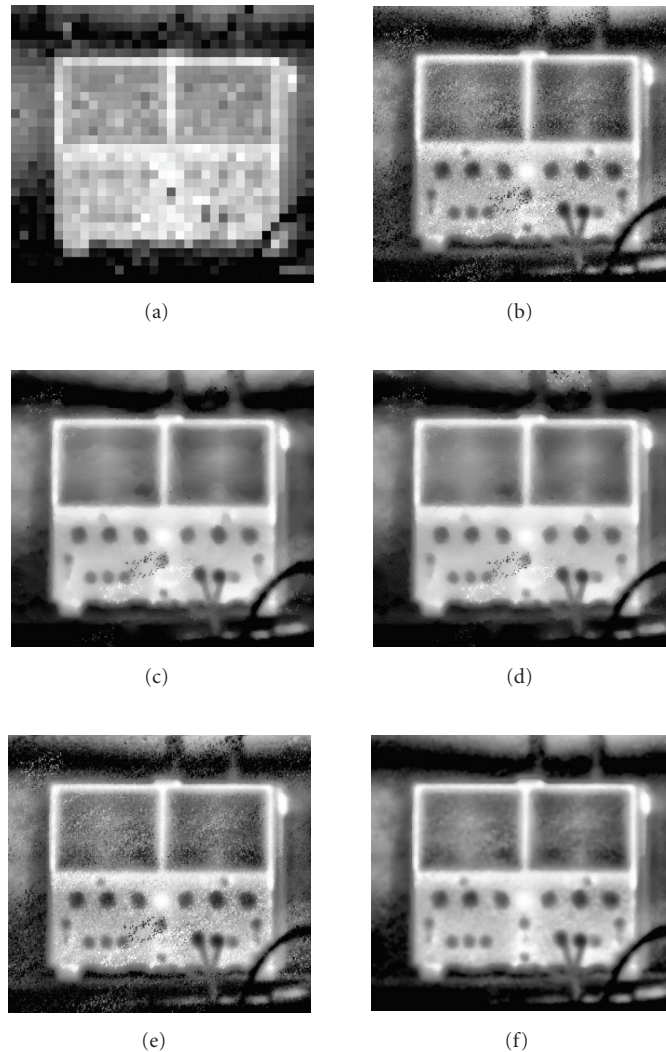
FIGURE 9: Eight-times SR results without deconvolution. All images are stretched using the same parameters [31]. (a) Pixel replication; (b) shift and add [8]; (c) Zomet $\sigma_{PSF} = 0$, $\lambda = 3 \times 10^{-4}$, $\beta = 5$; (d) Farsiu $\sigma_{PSF} = 0$, $\lambda = 0.0017$, $\beta = 5$; (e) cubic Delaunay; and (f) robust NC.

iterations of robust NC are performed, followed by one iterationof adaptive NC for highly oriented pixels (pixels whose anisotropy $A > 0.5$). Since the fill-factor is low, many details previously aliased in the LR images are now visible in the four-times HR image without the need of deconvolution. Due to a large degree of overdetermination (100 frames for 4×4 upsampling), noise is greatly reduced. Thanks to the robust component of the algorithm, the HR image also shows no trace of dead pixels, which appear abundantly in Figure 8(a) as highly dark and bright pixels.

To better visualize the capability of robust NC, we perform eight-times SR of a small region of interest (ROI) and show the results in Figure 9. The ROI renders an apparatus with many small features of various sizes that are useful for visual inspection. Images in the top row are a LR image and a nonrobust fusion results using a quick shift and add (S&A) method [8]. As can be seen in Figure 9(b), the S&A image is no longer aliased as the LR input and many small details are clearly visible. This substantial improvement in resolution is

a direct result of accurate motion vectors computed by the optimal shift estimator [18]. According to the performance limit finding in [18], these motion estimates are accurate enough for an eight-times SR because the motion is computed over big and high SNR images.

However, being a nonrobust fusion method, S&A cannot reduce noise and outliers from a low number of samples setting (100 frames for an 8×8 upsampling). Because the S&A result is often used as an initialization to the Zomet and Farsiu methods [9], these methods also suffer from the outliers left behind by S&A. The effect can clearly be seen in the visually best fusion results of Zomet and Farsiu in the middle row of Figure 9. These images are produced without a deconvolution kernel to be comparable with other fusion-only methods in Figure 9. Although designed to be robust, these two methods can remove low noise but not strong outliers (very dark or very bright pixels in the S&A image). The use of a higher regularization parameter λ does not improve the situation either, because small details in the image start to

(a)   (b)

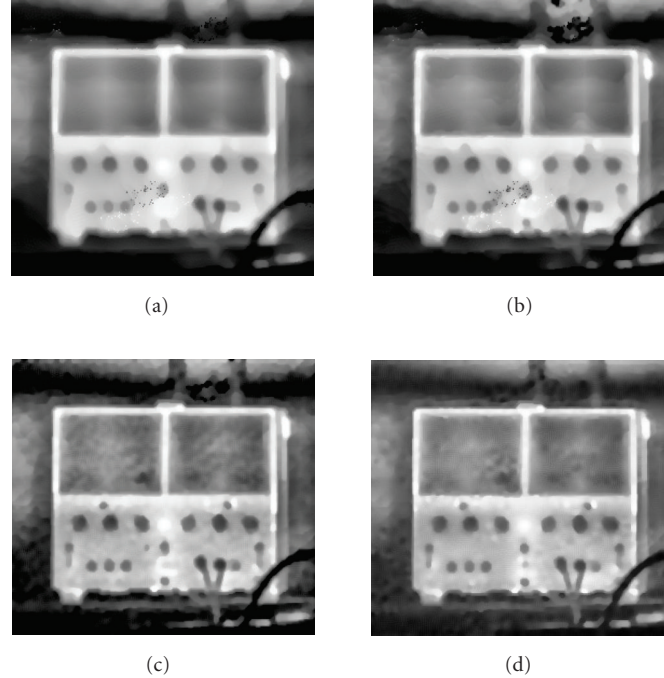(c)   (d)

FIGURE 10: Results of 8-time SR with bilateral TV deconvolution. All images are stretched using the same parameters [31]. (a) Zomet + bilateral TV regularization ($\lambda = 0.002$, $\beta = 2$); (b) Farsiu S&A followed by L2 + bilateral TV regularization ($\lambda = 0.002$, $\beta = 2$); (c) S&A followed by L1 + bilateral TV deconvolution ($\lambda = 0.1$, $\beta = 8$); and (d) robust NC followed by L1 + bilateral TV deconvolution ($\lambda = 0.05$, $\beta = 20$).

dissolve as λ increases (e.g., the two small circles just below the two display panels of the apparatus are barely visible in Figures 9(c) and 9(d)).

The last row of Figure 9 shows the results of SR fusion from two surface interpolation methods: a nonrobust fusion method using Delaunay triangulation [15] and a robust local surface fit using adaptive NC. For this type of noisy data, a surface interpolator that goes through every data point performs no better than the fast and simple S&A method in Figure 9(b). In fact, noise is even enhanced in Figure 9(b) because piecewise cubic interpolation is applied to the Delaunay tessellation. On the contrary, the adaptive NC result shows a high level of details without any artifacts. This is the strongest point of adaptive NC over other presented methods (robust and nonrobust alike) because it properly preconditions the HR image for the final deconvolution step.

5.2. Super-resolution by deconvolution

While fusion achieves some resolution enhancement under the presence of aliasing, deconvolution is necessary to remove the blur caused by optics and sensor elements. In this subsection, we apply deconvolution to the fusion results in the previous subsection. The combined optics and sensor blur are considered to be Gaussian and the scale of this Gaussian PSF is found to be $\sigma_{PSF} = 2$ by fitting a Gaussian edge model to various step edges in the fusion image [16]. Since bilateral TV with an L2 data norm (L2 + bilateral TV) is incorporated in the Farsiu and Zomet implementations [9] prior to deconvlution, we show the visually best

results for these methods in Figures 10(a) and 10(b). However, we found that a norm-one data with bilateral TV prior deconvolution [9] (L1+ bilateral TV) performs better on this type of noisy IR data. Unfortunately, the software given by [9] does not incorporate L1 + bilateral TV deconvolution into the Zomet and Farsiu methods. As a result, we apply our own implementation of L1 + bilateral TV deconvolution to the S&A and adaptive NC fusion images and show the deblurred results in Figures 10(c) and 10(d).

The restoration results in the first row of Figure 10 show that Zomet and Farsiu methods still cannot remove the outliers from the S&A initialization. Although the Farsiu result performs slightly better than the Zomet result for the same set of parameters ($\sigma_{PSF} = 2$, $\lambda = 0.002$, $\beta = 2$), the difference is very subtle. The second variant of Farsiu method using L1 + bilateral TV deconvolution in Figure 10(c) produces a much better image than L2 + bilateral TV. However, since Figure 10(c) starts with a nonrobust S&A image, some outliers are not completely removed. More dangerously, spurious details created from those outliers can be mistakenly recognized as real details. For example, on the left of a real knob in the middle of the control panel appears a small dot that looks just like a tiny mark. Also, in the place of an outlier clutter on top of image, there are now stain marks as a result of TV regularization. The deblurred NC result in Figure 10(d) shows none of these disturbing artifacts. Moreover, very fine details are resolvable like a real dot just below the same knob in the middle. This small dot is almost invisible in the S&A and NC images in Figures 9(b) and 9(f), and it only becomes clear in Figure 10(d) after an

L1 + bilateral TV deconvolution. In short, the robust and adaptive NC is preferable over the nonrobust S&A fusion method. This is especially true when fusion images undergo deconvolution because low input noise requires less regularization, which in turns improves detail restoration.

6. CONCLUSIONS AND DISCUSSIONS

We propose a solution for fusion of irregularly sampled images using adaptive normalized convolution. The method performs a robust polynomial fit over an adaptive neighborhood. Each sample could carry its own certainty or is automatically assigned a robust certainty based on the intensity difference against the central pixel in the current analysis window. The novelty of the method lies in the adaptive applicability which extends along local orientation to gather more samples of the same modality for a better analysis. The applicability function also contracts in the normal orientation to prevent smoothing across lines and edges. The principle can be extended to curved anisotropic applicability functions using recent curvature estimation techniques [21, 22]. In addition, the robust sample certainty minimizes the smoothing of sharp corners and tiny details because samples from other intensity distributions are effectively ignored in the local analysis.

The effectiveness of robust fusion using adaptive NC has been demonstrated through the application of super-resolution reconstruction of LR image sequences. In SR fusion, adaptive NC outperforms other methods such as the Delaunay triangulation-based interpolation algorithm [15] and many iterative algorithms including regularized back-projection [12], robust fusion using median of back-projected errors [30], and robust fusion using bilateral total variation regularization [9]. Apart from producing a more detailed image reconstruction, adaptive NC fusion is also fast and robust against noise and outliers. Although the adaptive NC is presented for fusion of shifted image sequences, the algorithm is applicable to any problem of fusion of irregularly sampled signals.

Not only useful in fusion of irregularly sampled images, adaptive normalized convolution is also applicable to a number of other problems. In [19], we use zero-order adaptive NC to perform geometry-driven image inpainting. The adaptive applicability function can be integrated into many other techniques including bilateral filtering for edge-preserving smoothing [24], robust Gaussian facet model for orientation estimation [27], and polynomial expansion for motion estimation [7]. Finally, the robust signal certainty presented in this paper can be utilized in some non-interpolating fusion technique such as thin-plate spline interpolation [10] to reduce the influence of outliers.

ACKNOWLEDGMENT

REFERENCES

[1] I. Amidror, "Scattered data interpolation methods for electronic imaging systems: a survey," *Journal of Electronic Imaging*, vol. 11, no. 2, pp. 157–176, 2002.

[2] S. Borman and R. L. Stevenson, "Super-resolution from image sequences: a review," in *Proceedings of Midwest Symposium on Circuits and Systems (MWSCAS '98)*, pp. 374–378, Notre Dame, Ind, USA, August 1998.

[3] D. R. K. Brownrigg, "The weighted median filter," *Communications of the ACM*, vol. 27, no. 8, pp. 807–818, 1984.

[4] D. Capel, *Image Mosaicing and Super-Resolution*, Springer, Berlin, Germany, 2004.

[5] M. Elad, "On the origin of the bilateral filter and ways to improve it," *IEEE Transactions on Image Processing*, vol. 11, no. 10, pp. 1141–1151, 2002.

[6] M. Elad and Y. Hel-Or, "A fast super-resolution reconstruction algorithm for pure translational motion and common space-invariant blur," *IEEE Transactions on Image Processing*, vol. 10, no. 8, pp. 1187–1193, 2001.

[7] G. Farnebäck, *Polynomial expansion for orientation and motion estimation*, Ph.D. thesis, Linköping University, Linköping, Sweden, 2002.

[8] S. Farsiu, D. Robinson, M. Elad, and P. Milanfar, "Robust shift and add approach to superresolution," in *Applications of Digital Image Processing XXVI*, vol. 5203 of *Proceedings of SPIE*, pp. 121–130, San Diego, Calif, USA, August 2003.

[9] S. Farsiu, M. D. Robinson, M. Elad, and P. Milanfar, "Fast and robust multiframe super resolution," *IEEE Transactions on Image ProcessinG*, vol. 13, no. 10, pp. 1327–1344, 2004.

[10] R. Franke, "Smooth interpolation of scattered data by local thin plate splines," *Computers & Mathematics with Applications*, vol. 8, no. 4, pp. 273–281, 1982.

[11] R. M. Haralick and L. Watson, "A facet model for image data," *Computer Graphics and Image Processing*, vol. 15, no. 2, pp. 113–129, 1981.

[12] 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," *Optical Engineering*, vol. 37, no. 1, pp. 247–260, 1998.

[13] M. Irani and S. Peleg, "Improving resolution by image registration," *CVGIP: Graphical Models and Image Processing*, vol. 53, no. 3, pp. 231–239, 1991.

[14] H. Knutsson and C.-F. Westin, "Normalized and differential convolution," in *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR '93)*, pp. 515–523, New York, NY, USA, June 1993.

[15] S. Lertrattanapanich and N. K. Bose, "High resolution image formation from low resolution frames using Delaunay triangulation," *IEEE Transactions on Image Processing*, vol. 11, no. 12, pp. 1427–1441, 2002.

[16] M. Luxen and W. Förstner, "Characterizing image quality: blind estimation of the point spread function from a single image," in *Proceedings of Photogrammetric Computer Vision (PCV '02)*, pp. 205–210, Graz, Austria, September 2002.

[17] M. Nitzberg and T. Shiota, "Nonlinear image filtering with edge and corner enhancement," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 8, pp. 826–833, 1992.

[18] T. Q. Pham, M. Bezuijen, L. J. van Vliet, K. Schutte, and C. L. Luengo Hendriks, "Performance of optimal registration

estimators," in *SPIE Defense and Security Symposium, Visual Information Processing XIV*, vol. 5817 of *Proceedings of SPIE*, pp. 133–144, Orlando, Fla, USA, March–April 2005.

[19] T. Q. Pham and L. J. van Vliet, "Normalized averaging using adaptive applicability functions with applications in image reconstruction from sparsely and randomly sampled data," in *Proceedings of 13th Scandinavian Conference on Image Analysis (SCIA '03)*, vol. 2749 of *Lecture Notes in Computer Science*, pp. 485–492, Göteborg, Sweden, June–July 2003.

[20] T. Q. Pham, L. J. van Vliet, and K. Schutte, "Influence of signal-to-noise ratio and point spread function on limits of superresolution," in *IS&T/SPIE's 17th Annual Symposium Electronic Imaging Science and Technology, Image Processing: Algorithms and Systems IV*, vol. 5672 of *Proceedings of SPIE*, pp. 169–180, San Jose, Calif, USA, January 2005.

[21] B. Rieger, F. J. Timmermans, L. J. van Vliet, and P. W. Verbeek, "On curvature estimation of ISO surfaces in 3D gray-value images and the computation of shape descriptors," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 8, pp. 1088–1094, 2004.

[22] B. Rieger and L. J. van Vliet, "Curvature of n-dimensional space curves in grey-value images," *IEEE Transactions on Image Processing*, vol. 11, no. 7, pp. 738–745, 2002.

[23] K. Schutte, D.-J. J. de Lange, and S. P. van den Broek, "Signal conditioning algorithms for enhanced tactical sensor imagery," in *Infrared Imaging Systems: Design, Analysis, Modeling, and Testing XIV*, vol. 5076 of *Proceedings of SPIE*, pp. 92–100, Orlando, Fla, USA, April 2003.

[24] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," in *Proceedings of 6th International Conference on Computer Vision (ICCV '98)*, pp. 839–846, Bombay, India, January 1998.

[25] J. van de Weijer and R. van den Boomgaard, "Local mode filtering," in *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR '01)*, vol. 2, pp. 428–433, Kauai, Hawaii, USA, December 2001.

[26] J. van de Weijer, L. J. van Vliet, P. W. Verbeek, and M. van Ginkel, "Curvature estimation in oriented patterns using curvilinear models applied to gradient vector fields," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23, no. 9, pp. 1035–1042, 2001.

[27] R. van den Boomgaard and J. van de Weijer, "Linear and robust estimation of local image structure," in *Proceedings of 4th International Conference on Scale-Space Theories in Computer Vision (Scale-Space '03)*, vol. 2695 of *Lecture Notes in Computer Science*, pp. 237–254, Isle of Skye, Scotland, UK, June 2003.

[28] R. A. Young, "The Gaussian derivative model for spatial vision: I. Retinal mechanisms," *Spatial Vision*, vol. 2, no. 4, pp. 273–293, 1987.

[29] I. T. Young, L. J. van Vliet, and M. van Ginkel, "Recursive Gabor filtering," *IEEE Transactions on Signal Processing*, vol. 50, no. 11, pp. 2798–2805, 2002.

[30] A. Zomet, A. Rav-Acha, and S. Peleg, "Robust super-resolution," in *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR '01)*, vol. 1, pp. 645–650, Kauai, Hawaii, USA, December 2001.

[31] K. Zuiderveld, "Contrast limited adaptive histogram equalization," in *Graphics Gems IV*, P. S. Heckbert, Ed., pp. 474–485, Academic Press, Boston, Mass, USA, 1994.

Tuan Q. Pham was born in Vietnam in 1978. In 1997 he won an AusAID scholarship to study in Monash University, Australia, where he obtained his Bachelor of Computer Science and Engineering with first class honors (2001). In 2002, he joined the Pattern Recognition Group at the Delft University of Technology, The Netherlands, to commence his Ph.D. research on "Super-resolution of under-sampled image sequences." His current research interests include structure adaptive filtering, fusion of uncertain and irregularly sampled signals, motion estimation, and super-resolution. He is a silver medallist at the 36th International Mathematical Olympiad held in Canada, 1995.
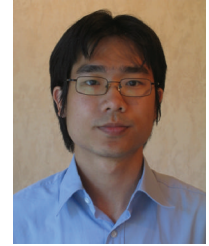
Lucas J. Van Vliet (1965) is a Full Professor in multidimensional data analysis at the Faculty of Applied Sciences of the Delft University of Technology in The Netherlands. He received his M.S. degree in applied physics in 1988 and his Ph.D. degree cum laude in 1993. His thesis entitled "Grey-scale measurements in multidimensional digitized images" presents novel methods for sampling-error-free measurements of geometric object features. He has worked on various sensor, restoration, and measurement problems in quantitative microscopy. His current research interests include segmentation and analysis of objects, textures and structures in multidimensional digitized images from a variety of imaging modalities. In 1996 he was awarded a fellowship of the Royal Netherlands Academy of Arts and Sciences (KNAW).

Klamer Schutte performed Ph.D. work at the University Twente, and graduated in 1994 on his thesis "Knowledge based recognition of man-made objects." After a two-year stay as Post-Doc at the Pattern Recognition Group of the Delft University of Technology, he joined TNO Physics and Electronics Laboratory in 1996. His current position is Chief Scientist in Electro-optics Group of TNO Defence, Security, and Safety.

EUSIPCO 2011

19th European Signal Processing Conference
August 29- September 2, 2011, Barcelona (Spain)

Photograph © Turisme de Barcelona / J. Trullàs

Preliminary call for papers

The 2011 European Signal Processing Conference (EUSIPCO-2011) is the nineteenth in a series of conferences promoted by the European Association for Signal Processing (EURASIP, www.eurasip.org). This year edition will take place in Barcelona, capital city of Catalonia (Spain), and will be jointly organized by the Centre Tecnològic de Telecomunicacions de Catalunya (CTTC) and the Universitat Politècnica de Catalunya (UPC).

EUSIPCO-2011 will focus on key aspects of signal processing theory and applications as listed below. Acceptance of submissions will be based on quality, relevance and originality. Accepted papers will be published in the EUSIPCO proceedings and presented during the conference. Paper submissions, proposals for tutorials and proposals for special sessions are invited in, but not limited to, the following areas of interest.

Areas of Interest

- Audio and electro-acoustics.
- Design, implementation, and applications of signal processing systems.
- Multimedia signal processing and coding.
- Image and multidimensional signal processing.
- Signal detection and estimation.
- Sensor array and multi-channel signal processing.
- Sensor fusion in networked systems.
- Signal processing for communications.
- Medical imaging and image analysis.
- Non-stationary, non-linear and non-Gaussian signal processing.

Submissions

Procedures to submit a paper and proposals for special sessions and tutorials will be detailed at www.eusipco2011.org. Submitted papers must be camera-ready, no more than 5 pages long, and conforming to the standard specified on the EUSIPCO 2011 web site. First authors who are registered students can participate in the best student paper competition.

Important Deadlines:

Proposals for special sessions	15 Dec 2010
Proposals for tutorials	18 Feb 2011
Electronic submission of full papers	**21 Feb 2011**
Notification of acceptance	23 May 2011
Submission of camera-ready papers	6 Jun 2011

Webpage: www.eusipco2011.org

Organizing Committee

Honorary Chair
Miguel A. Lagunas (CTTC)

General Chair
Ana I. Pérez-Neira (UPC)

General Vice-Chair
Carles Antón-Haro (CTTC)

Technical Program Chair
Xavier Mestre (CTTC)

Technical Program Co-Chairs
Javier Hernando (UPC)
Montserrat Pardàs (UPC)

Plenary Talks
Ferran Marqués (UPC)
Yonina Eldar (Technion)

Special Sessions
Ignacio Santamaría (Unversidad de Cantabria)
Mats Bengtsson (KTH)

Finances
Montserrat Nájar (UPC)

Tutorials
Daniel P. Palomar
(Hong Kong UST)
Beatrice Pesquet-Popescu (ENST)

Publicity
Stephan Pfletschinger (CTTC)
Mònica Navarro (CTTC)

Publications
Antonio Pascual (UPC)
Carles Fernández (CTTC)

Industrial Liaison & Exhibits
Angeliki Alexiou
(University of Piraeus)
Albert Sitjà (CTTC)

International Liaison
Ju Liu (Shandong University-China)
Jinhong Yuan (UNSW-Australia)
Tamas Sziranyi (SZTAKI -Hungary)
Rich Stern (CMU-USA)
Ricardo L. de Queiroz (UNB-Brazil)