

Robust wavelet-based super-resolution reconstruction: Theory and Algorithm

Hui Ji, Cornelia Fermüller *Member, IEEE*

Abstract

We present an analysis and algorithm for the problem of super-resolution imaging, that is the reconstruction of HR (high-resolution) images from a sequence of LR (low-resolution) images. Super-resolution reconstruction entails solutions to two problems. One is the alignment of image frames. The other is the reconstruction of a HR image from multiple aligned LR images. Both are important for the performance of super-resolution imaging. Image alignment is addressed with a new batch algorithm, which simultaneously estimates the homographies between multiple image frames by enforcing the surface normal vectors to be the same. This approach can handle video sequences with large displacement quite well. Reconstruction is addressed with a wavelet-based iterative reconstruction algorithm with an efficient de-noising scheme. The technique is based on a new analysis of video formation. At a high level our method could be described as a better-conditioned iterative back projection scheme with an efficient regularization criteria in each iteration step. Experiments with both simulated and real data demonstrate that our approach has better performance than existing super-resolution methods. It can remove even large amounts of mixed noise without creating artifacts.

Index Terms

Perfect reconstruction filter banks, Super-resolution, Multiple frame alignment, Wavelet denoising

I. OVERVIEW

The problem of *super-resolution* reconstruction, which is defined as restoring a high-resolution (HR) image from a sequence of low-resolution images (see Fig. I), has been studied by many researchers in recent years. Most super-resolution algorithms formulate the problem as a signal

Hui Ji is with the Department of Mathematics, National University of Singapore, Singapore

Cornelia Fermüller is with the Center for Automation Research, University of Maryland, College Park, USA

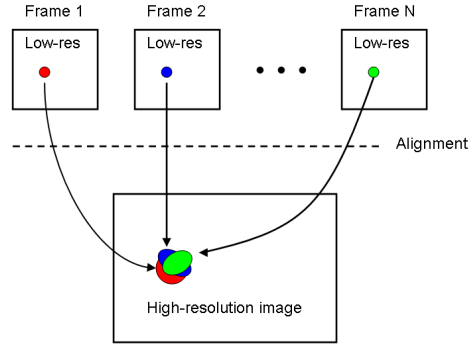


Fig. 1. Illustration of super-resolution imaging.

reconstruction problem. Essentially these algorithms differ in two aspects: one is in how the image frames of the sequence are aligned in a common coordinates system; the other is in how the high-resolution image is reconstructed from the aligned low-resolution image frames. The accuracy of the image alignment and the robustness of the signal reconstruction determine the algorithmic performance. The two aspects are related. Thus the imperfectness of the image alignment has to be considered in the later stage of image restoration.

In this paper we propose a complete super-resolution reconstruction system which addresses both the problems of frame alignment and image restoration. Our contributions are two fold. The first is a theoretical analysis of video formation from the view of filter banks, which leads to a new wavelet-based reconstruction scheme for the image reconstruction. The second is a batch algorithm for image alignment, which computes the projective flow (planar homographies) across all frames in the sequences.

Using a sequence of similar but not identical image frames makes it is possible to increase the spatial resolution of one still image. However, it is difficult since this is an ill-posed inverse problem. Despite the difficulties, researchers have made great progress towards stable algorithms. Iterative back-projection methods ([10], [18]) have been shown to be an effective tool for high-resolution image reconstruction. It is known, however, that the de-blurring process, which is part of this approach, makes it very sensitive to the noise. Thus, the requirement of very accurate image alignment estimates limits its practical use. Various regularization methods have been proposed to deal with the noise. However, these methods either are very sensitive to the

assumed noise model (Tikhonov regularization) or are computationally expensive (Total-Variation regularization). See [6] for more details.

Our contributions to the reconstruction process are two-fold. First, we model the image formation procedure from the point of view of filter bank theory. Then, based on this new formulation, we analyze the limits of the high-resolution reconstruction. Our conclusion is that in general full recovery is not possible without enforcing some constraints on the recovered images. At best we could reconstruct the image convolved with a specific low-pass filter (namely $\frac{1}{4}(1, 1) \otimes (1, 1)$ for the case of the Box-type PSF). Second, based on our new formulation, we present a robust wavelet-based algorithm for image reconstruction. The iteration scheme in our algorithm is inherently more robust to noise than that of classic back-projection methods ([10], [18]), because the projection matrix of our back-projection scheme has a better condition number. We will show that, both in theory and experiments, it has better performance in suppressing the error propagation than other back-projection iteration schemes.

Furthermore, our algorithm allows us to include a wavelet-based de-noising scheme in each iteration of the reconstruction which effectively removes the noise without creating smoothing artifacts. The advantage of our de-noising scheme over regularization methods is that it is nearly optimal with respect to the risk bound. That is, it has the theoretical minimal error in removing noise of unknown models. Its effectiveness in removing mixed noise and relatively large amounts of noise is demonstrated in experiments. It is worth mentioning that our de-noising scheme adds very little computational burden compared to other complicated regularization methods. Briefly, our method could be described as a generalized iterative back-projection method with a fast and optimal regularization criteria in each iteration step. Wavelet theory has previously been used for image de-noising and de-blurring from static images ([3], [4]). However, it has not been studied much with respect to the super-resolution problem. In recent work wavelet theory has been applied to this problem [16], but only for the purpose of speeding up the computation. Our contribution lies in an analysis that reveals the relationship between the inherent structure of super-resolution reconstruction and the theory of wavelet filter banks. This relationship is fully exploited by using various techniques from wavelet theory in the iterations of the reconstruction.

For the process of image frame alignment, the flow-based approach is the most popular choice because of its flexibility. The difficulty in accurate estimation of general non-parametric image flow makes the reconstruction of higher-resolution images not very meaningful. Thus, some

assumption on the underlying flow model has to be enforced. Such assumption, of course, makes the algorithm less general, but more practical. The most frequently used flow models are the affine transform and the more general projective transform ([2], [5], [10], [22]). The most popular reconstruction procedure is the iterative back projection method ([10], [18], [23]).

In this paper, we propose a new algorithm for computing projective flow (planar homographies) simultaneously on all frames in the sequence. The most closely related studies are [8], [21] and [11], which all provide algorithms for the estimation of planar motion between multiple frames based on some rank constraint. However these studies either are limited to short image sequences ([8], [21]), or compute relative motion only between neighboring frames ([11]), which could lead to accumulative errors in the image alignment. Our approach is an extension of [11] to the more general estimation of homographies (as opposed to planar flow). Briefly, the motion between the reference frame and any other frame is modeled as a homography. Then, the shape constraint proposed in [11] is directly incorporated into the estimation of the homographies. This leads to accurate alignment for longer sequences. As a result, we have more frames available for the reconstruction.

II. FORMULATION OF HIGH-TO-LOW IMAGE FORMATION

We first formulate the high-to-low image formation process. To simplify the exposition, in the following we only discuss 1D signals with resolution enhancement by a factor 2. Later, without much difficulty, the analysis will be extended to the 2D case with arbitrary resolution enhancement. Adopting Farsiu's notation ([6]), the image formation process in the pixel domain can be modeled as

$$y = \sigma[H * X(F(t))] + N, \quad (1)$$

where t is the spatial variable; $X(t)$ is the continuous signal and y is the discrete signal; H is the blurring operator (either optical blurring or motion blurring or both); F is the geometric transform; N is the noise in the low-resolution image; σ is the decimation operator; and “*” is the convolution operator. Not considering the noise for the moment, the high-resolution (HR) signal x and the low-resolution (LR) signal y can be defined as:

$$x = \sigma[X], \quad y = [\sigma[H * X(F(t))]] \downarrow_2, \quad (2)$$

where \downarrow_2 is the downsampling operator with rate 2.

Next we derive the relationship between the LR signal y and the HR signal x . Denote the velocity of the signal as $\epsilon(t) = F(t) - t$, which is also called the *optical flow* in Computer Vision. For simplicity of notation, here we assume a sub-pixel flow model with $0 < \epsilon(t) < 2$ on the denser grid of the HR image x (Larger flow could always be reduced to the case of sub-pixel flow by re-assigning the pixel value). Thus, in the LR image the flow values are all sub-pixel shifts (Recall a 1-unit shift on the coarse grid of y equals a 2-unit shift on the fine grid of x).

Let $\{j\}$ be a fine grid for the spatial coordinates x . Then for point j of y on the coarse grid (its coordinate is $2j$ on the fine grid) with $0 \leq \epsilon(2j) < 1$, the first-order Taylor approximation of Equation (2) at point $2j$ can be written as

$$\begin{aligned} y(j) &= [H * X(F(t))]_{t=2j} \\ &= [H * X(\epsilon(t) + t)]_{t=2j} \\ &= [H * X(t)]_{t=2j} + \epsilon(2j)[H * X'(t)]_{t=2j} \\ &= [H * X(t)]_{t=2j} + \epsilon(2j)[H' * X(t)]_{t=2j}. \end{aligned}$$

For all other points j' of y with $1 \leq \epsilon(2j') < 2$, a similar argument yields

$$y(j') = [H * X(t)]_{t=2j'+1} + (\epsilon(2j') - 1)[H' * X(t)]_{t=2j'+1}.$$

Thus, a LR sequence y could be expressed in the pixel domain as a sub-sequence of the following two sequences:

$$[a * x] \downarrow_2 + \epsilon \cdot [b * x] \downarrow_2; \quad \text{and} \quad [a * x(\cdot + 1)] \downarrow_2 + (\epsilon - 1) \cdot [b * x(\cdot + 1)] \downarrow_2, \quad (3)$$

where a, b are discrete versions of the convolution kernels H and H' respectively, and $\cdot*$ denotes the component-wise multiplication operator. Having available the optical flow values ϵ_k for multiple low-resolution images y_k , we can by solving an over-determined system of linear equations obtain the four components:

$$[a * x] \downarrow_2, \quad [a * x(\cdot + 1)] \downarrow_2, \quad [b * x] \downarrow_2, \quad [b * x(\cdot + 1)] \downarrow_2. \quad (4)$$

As will be shown in the next subsection, the two filters a and b (which are determined by the blurring kernel H and its derivative H') characterize the super-resolution reconstruction.

Let us next look at some examples of filters a and b for different blurring kernels.

Example 1 Consider the box-type blurring kernel $H = \frac{1}{2n}\chi_{[-n,n]}$. Let $\epsilon(t) = \epsilon \leq 1$. Then we have

$$\begin{aligned} y(j) &= \int_{-\infty}^{\infty} \chi(2j-t)X(F(t))dt \\ &= \frac{1}{2n} \int_{2j-n}^{2j+n} X(F(t))dt = \frac{1}{2n} \int_{2j-n}^{2j+n} X(t+\epsilon)dt. \end{aligned}$$

Approximating the integration by quadrature rules, we obtain

$$y(j) = \frac{1}{2n} \left(\frac{1}{2}(1-\epsilon)x(2j-n) + \sum_{i=-n+1}^{n-1} x(2j-i) + \frac{1}{2}(1+\epsilon)x(2j+n) \right).$$

Equivalently, we can write

$$y = [a * x + \epsilon(b * x)] \downarrow_2, \quad (5)$$

where a and b are the following low-pass and high-pass filters, respectively:

$$a = \frac{1}{4n}(1, 2, \dots, 2, 1), \quad b = \frac{1}{4n}(-1, 0, \dots, 0, 1).$$

Example 2 Consider a Gaussian-type blurring kernel H . Using the Cubic Cardinal B-spline $B(t)$ as approximation to the Gaussian function we have

$$y(j) = \int_{-\infty}^{\infty} B(2j-t)X(F(t))dt.$$

Again, by the quadrature rule, we have the approximation

$$y = \sum_i x(2j-i)(a(i) - \epsilon \cdot b(i)),$$

where

$$\begin{cases} a = \frac{1}{96}(1, 8, 23, 32, 23, 8, 1); \\ b = \frac{1}{48}(3, 12, 15, 0, -15, -12, -3). \end{cases}$$

III. ANALYSIS OF THE HR RECONSTRUCTION

Given multiple LR signals y_k with different motions ϵ_k , theoretically we can obtain two complete sequences $a*x$ and $b*x$ from (4). An interesting question arises. Without any assumption on the given finite signal x , can we reconstruct the signal sequence x exactly from the two sequences $a * x$ and $b * x$?

To answer this question, let us write the sequence in another form, namely, as its Z-transform. The Z-transform of a signal sequence $x = \{x(i)\}$ is defined as

$$x(z) = \sum_i x(i)z^{-i}.$$

It is easy to see that this transform is a one-to-one mapping between sequence space and polynomial space. Let $a(z)$ and $b(z)$ denote the Z-transforms of the filters a and b , then the Z-transforms of $a * x$ and $b * x$ are $a(z)x(z)$ and $b(z)x(z)$ respectively.

Now the question can be addressed by checking whether the polynomial equation

$$(a(z)x(z))u(z) + (b(z)x(z))v(z) = x(z), \quad (6)$$

is solvable for the two unknowns $u(z)$ and $v(z)$. Eliminating $x(z)$ from both sides of (6) yields

$$a(z)u(z) + b(z)v(z) = 1. \quad (7)$$

From the theory of Diophantine equation we know the following:

Lemma 3 *Given two polynomials $a(z)$ and $b(z)$, (7) is solvable if and only if the greatest common divisor of $a(z)$ and $b(z)$ is a scalar, that is, if $a(z)$ and $b(z)$ are co-prime.*

It is observed that $a(z)$ and $b(z)$ in our two examples (Example 1 and Example 2) both have a common divisor

$$c(z) = (1 + z).$$

This can be seen from the fact that $a(-1) = b(-1) = 0$, and therefore $z = -1$ is the root of both $a(z)$ and $b(z)$. Thus, for these blurring kernels we cannot reconstruct $x(z)$ from $a(z)x(z)$ and $b(z)x(z)$ exactly. This observation is not an incident. The same holds true for general blurring kernels, as we will show next.

We follow Baker's modeling of the blurring kernel H ([1]). The blurring kernel (Point spread function) can be decomposed into two components:

$$H = \Omega * C,$$

where $\Omega(X)$ models the blurring caused by the optics and $C(X)$ models the spatial integration performed by the CCD sensor. Typically Ω is modeled by a Gaussian-type function and C is modeled by a Box-type function. Notice that

$$H' = \Omega' * C.$$

Thus we can express the corresponding discrete filters as:

$$a = \ell * c; \quad b = \tau * c,$$

where c is the discrete version of the spatial integration kernel C , and ℓ and τ are the discrete versions of Ω and Ω' . Since $a(z)$ and $b(z)$ have a common divisor $c(z)$, we cannot reconstruct $x(z)$ for general $x(z)$, unless C is a Dirac function, which generally is not true. Based on Lemma 3, we then have the following result.

Theorem 4 *Given multiple LR finite signals y_k , we can not perfectly reconstruct the HR finite signal x without any assumptions on x . At most we can reconstruct $c*x$ for some low-pass filter c . The corresponding Z-transform $c(z)$ of c is the greatest common divisor of $a(z)$ and $b(z)$, which includes the spatial integration filter.*

Notice that c is a low-pass FIR (finite impulse response) filter. To recover x from $c*x$, we have to apply a high-pass filter on $c*x$ and impose some boundary condition on the signal x . Such a de-blurring process generally is sensitive to the noise and creates artifacts in the recovered image. A good strategy then is to modify our reconstruction goal during the intermediate iterative reconstruction process: Instead of trying to reconstruct x , we reconstruct $c*x$ in the iterative process, and we leave the recovery of x from $c*x$ to the last step, after finishing the iterative reconstruction.

Thus, the modified HR signal to be reconstructed is $\tilde{x} = c*x$. The LR sequence $\{y_k\}$ can be expressed as a subset of the following two sequences:

$$[\ell * \tilde{x}] \downarrow_2 + \epsilon_k [\tau * \tilde{x}] \downarrow_2 \tag{8}$$

and

$$[\ell * \tilde{x}(\cdot + 1)] \downarrow_2 + (\epsilon_k - 1) [\tau * \tilde{x}(\cdot + 1)] \downarrow_2,$$

where the Z-transform of ℓ and τ are $a(z)$ and $b(z)$ divided by their greatest common divisor $c(z)$, respectively. It is worth mentioning that we model the blurring procedure from HR to LR by a first-order Taylor approximation. But our reasoning could easily be extended to a modeling by higher-order Taylor approximations, leading to the same conclusion.




Fig. 2. The two-channel filter bank.

IV. RECONSTRUCTION BASED ON PR FILTER BANKS

A. Introduction to PR filter banks

Before presenting our algorithm, we first give a brief introduction to 2-channel PR (*perfect reconstruction*) filter banks, also called wavelet filter banks (see [14] for more details). A two-channel filter bank consists of two parts: an analysis filter bank and a synthesis filter bank. In our case, the signal \tilde{x} is first convolved with a low-pass filter ℓ and a high-pass filter h and then subsampled by 2. In other words, we analyze the signal by an analysis filter bank. Then a reconstructed signal \hat{x} is obtained by upsampling the signal by zero interpolation and then filtering it with a dual low-pass filter g and a dual high-pass filter q . In other words, we reconstruct the signal by synthesizing the output from the analysis bank with a synthesis filter bank. See Fig. 2 for an illustration.

Such a filter bank is called a PR filter bank if $\hat{x} = \tilde{x}$ for any input \tilde{x} . The question then is, what makes $\{\ell, h, q, g\}$ a PR filter bank. It is easy to see that the process illustrated in Fig. 2 could be expressed using Z-transforms as follows:

$$\hat{x}(z) = \frac{1}{2}(\tilde{x}(z) + \tilde{x}(-z)) \begin{pmatrix} \ell(z) & h(z) \\ \ell(-z) & h(-z) \end{pmatrix} \begin{pmatrix} g(z) \\ q(z) \end{pmatrix}.$$

Thus, the sufficient and necessary condition for $\hat{x}(z) = \tilde{x}(z)$ is

$$\begin{pmatrix} 2z^m \\ 0 \end{pmatrix} = \begin{pmatrix} \ell(z) & h(z) \\ \ell(-z) & h(-z) \end{pmatrix} \begin{pmatrix} g(z) \\ q(z) \end{pmatrix},$$

or equivalently,

$$\begin{pmatrix} g(z) \\ q(z) \end{pmatrix} = \begin{pmatrix} \ell(z) & h(z) \\ \ell(-z) & h(-z) \end{pmatrix}^{-1} \begin{pmatrix} 2z^m \\ 0 \end{pmatrix}.$$

Then in order to obtain FIR filters g and q , we require that

$$\det(H(z)) = \det\left(\begin{pmatrix} \ell(z) & h(z) \\ \ell(-z) & h(-z) \end{pmatrix}\right) = cz^n.$$

In summary, the analysis filters $\{\ell, h\}$ of a perfect reconstruction filter bank have to satisfy the following condition:

$$\ell(z)h(-z) - \ell(-z)h(z) = z^m \quad \text{for some integer } m, \quad (9)$$

and the corresponding synthesis filters amount to

$$g(z) = h(-z); \quad q(z) = -(\ell(-z)).$$

Thus, given a low-pass filter $\ell(z)$, we can find the corresponding high-pass filter $h(z)$ such that we have an analysis filter bank and thus a PR filter bank by solving the linear system (9).

Example 5 For the well-known ‘‘Harr’’ wavelet filter bank, the synthesis and analysis filters amount to:

$$\begin{aligned} \ell &= \frac{1}{2}(1, 1), & h &= \frac{1}{2}(1, -1); \\ g &= \frac{1}{2}(1, 1), & q &= \frac{1}{2}(-1, 1). \end{aligned}$$

B. Iterative reconstruction scheme

We have available a number of signals y_k and the corresponding estimates of the optic flow values ϵ_k . We also have estimates of the convolution kernels ℓ and τ . Obviously, estimating $\ell * x$ and $\tau * x$ directly from Equation (3) would not be wise. For reasons of numerical stability special attention is necessary. Fortunately, the scheme of PR filter banks provides us with an iterative scheme.

Let ℓ which corresponds to the blurring kernel be the low pass filter of a PR filter bank. Then the corresponding high-pass filter h is computed by solving (9). Note h may be different from τ .

Recall that for each LR signal y_k , we have

$$y_k = [\ell * \tilde{x}] \downarrow_2 + \epsilon_k \cdot * [\tau * \tilde{x}] \downarrow_2.$$

(If the original signal is incomplete, this can be overcome by a simple interpolation.) From the above equation $[\ell * \tilde{x}] \downarrow_2$ is obtained as

$$[\ell * \tilde{x}] \downarrow_2 = y_k - \epsilon_k \cdot * [\tau * \tilde{x}] \downarrow_2. \quad (10)$$

The process of a signal \tilde{x} passing through a PR filter bank as shown in Fig. 2 can be expressed as:

$$\tilde{x} = g * [(\ell * \tilde{x}) \downarrow_2] \uparrow_2 + q * [(h * \tilde{x}) \downarrow_2] \uparrow_2. \quad (11)$$

Combining (10) and (11), we obtain the iterative reconstruction of \tilde{x} from K LR signals y_k as follows: At step $n + 1$

$$\tilde{x}^{n+1} = g * [(h * \tilde{x}^n) \downarrow_2] \uparrow_2 + g * \left(\frac{1}{K} \sum_{k=1}^K [y_k - \epsilon_k \cdot * (\tau * \tilde{x}^n) \downarrow_2] \uparrow_2 \right). \quad (12)$$

Theorem 6 *The iteration of Equation (12) converges to the true value x under the condition that*

$$\|g * \tau\| \leq \frac{1}{2}. \quad (13)$$

Proof: As in [10] for simplicity we omit the down-sampling and upsampling process, as well as the fusion process. That is, we write

$$y = \ell * \tilde{x} + \epsilon \cdot * (\tau * \tilde{x}).$$

Then the iteration amounts to

$$\tilde{x}^{(n+1)} = g * (y - \epsilon \cdot * (\tau * \tilde{x}^{(n)})) + q * (\tau * \tilde{x}^{(n)}). \quad (14)$$

Subtracting \tilde{x} on both sides of (14) yields

$$\begin{aligned} \tilde{x}^{(n+1)} - \tilde{x} &= g * (y - \epsilon \cdot * (\tau * \tilde{x}^{(n)})) + q * (\tau * \tilde{x}^{(n)}) - \tilde{x} \\ &= g * (\ell * \tilde{x} + \epsilon \cdot * (\tau * (\tilde{x} - \tilde{x}^{(n)}))) + q * (\tau * \tilde{x}^{(n)}) - \tilde{x}. \end{aligned}$$

Recall that we have

$$\tilde{x} = g * (\ell * \tilde{x}) + q * (h * \tilde{x}).$$

Then

$$\begin{aligned} \tilde{x}^{(n+1)} - \tilde{x} &= -g * (\epsilon \cdot * (\tau * (\tilde{x}^{(n)} - \tilde{x}))) + q * h * (\tilde{x}^{(n)} - \tilde{x}) \\ &= (g * (-\epsilon \cdot * \tau) + q * h) * (\tilde{x}^{(n)} - \tilde{x}). \end{aligned}$$

Let A denote the operator which represents

$$g * (-\epsilon \cdot * \tau * + q * h *).$$

Then the above equation can be rewritten as

$$\tilde{x}^{(n+1)} - \tilde{x} = A(\tilde{x}^{(n)} - x).$$

Since we have (See [14] for more details):

$$\|q * h\| = \|g * \ell\| = \frac{1}{2},$$

from the fact that $\|\epsilon\|_\infty < 1$, we obtain

$$\|A\| \leq \|g * \tau\| + \|q * h\| < \|g * \tau\| + \frac{1}{2}.$$

Thus $\|g * \tau\| \leq \frac{1}{2}$ is sufficient for the convergence of the iteration. ■

C. Relationship to other back-projection methods

Applying (11), we can rewrite (12) in the form

$$\begin{aligned} \tilde{x}^{n+1} &= (\tilde{x}^n - g * [\ell * (\tilde{x}^n) \downarrow_2] \uparrow_2 + g * \left(\frac{1}{K} \sum_{k=1}^K [y_k - \epsilon_k \cdot * (\tau * \tilde{x}^n) \downarrow_2] \uparrow_2 \right)) \\ &= \tilde{x}^n + g * \left(\frac{1}{K} \sum_{k=1}^K [y_k - (\ell * \tilde{x}^n \epsilon_k \cdot * (\tau * \tilde{x}^n)) \downarrow_2] \uparrow_2 \right). \end{aligned}$$

From this we can see that the iteration scheme presented here falls in the class of back-projection methods. But it has advantages over the usual back-projection iterations. Consider the well-known method by Irani and Peleg [10]. Its iteration can be described as:

$$x^{n+1} = x^n + \frac{1}{K} \sum_{k=1}^K T_k^{-1} \left(((y_k - [\ell * T_k(x^n)] \downarrow_2) \uparrow_2) * p \right), \quad (15)$$

where T_k is the geometric transform between y_k and \tilde{x} , and the high-pass filter p is the de-blurring kernel. Notice that the two methods differ in the de-blurring kernel: one uses g with $g(z) = h(-z)$ defined in (9); the other is p in (15), the approximate inverse filter of ℓ .

The requirement on p in (15) is

$$\|\delta - \ell * p\| < 1, \quad (16)$$

where δ is the ideal unit impulse response filter. In other words, p should be a good approximation for the inverse of ℓ . In comparison, g in our iteration only needs to be a companion filter for

the smooth filter ℓ with sufficient decay, such that condition (13) holds. This difference makes g more desirable than p . Let's investigate this in more detail.

The noise will be propagated exponentially as $O(\|p\|^n)$ in (15) and as $O(\|g\|^n)$ in (12). Generally the flexibility of $g(z)$ makes it possible to design a g that has much smaller norm than p . This leads to much better resistance to noise propagation. Here is an example.

Example 7 Consider $\ell = \frac{1}{4}(1, 2, 1)$. Then

$$g = (-1/8, -1/4, 3/4, -1/4, -1/8)$$

is a dual PR filter for ℓ with

$$\ell(z)g(-z) = -1 + 9z^{-2} + 16z^{-3} + 9z^{-4} - z^{-6}.$$

It is easy to check that $\|g\|_2$ is around 0.85. The minimum for the norm of all filters with the same length as g is around 1.1. The corresponding p is

$$p = \left(\frac{1}{2}, -\frac{2}{3}, \frac{4}{3}, -\frac{2}{3}, \frac{1}{2}\right).$$

In order to make the norm of p close to the norm of g , a lengthy p with much slower decay is necessary. Such a filter is not desirable since it causes artifacts, like the ‘‘ring’’ effect. This clearly indicates that our iteration scheme is more robust to noise and causes less artifacts in the reconstructed image.

V. ROBUST ALGORITHM ON 2D IMAGES WITH DE-NOISING

Next we generalize the algorithm to 2D images. Then we introduce a de-noising process during the iterative reconstruction to suppress the noise in the optical flow estimation. Furthermore, the algorithm is adjusted to handle outliers.

A. Extension to 2D images with a built-in de-noising process

All the previous analysis can be generalized using the tensor product. By an argument similar to the 1D case, we approximate the LR image I^{LR} with the HR image I^{HR} as follows:

$$I^{LR} = [(a \otimes a) * I^{HR} + u \cdot ((a \otimes b) * I^{HR}) + v \cdot ((b \otimes a) * I^{HR})] \downarrow_2,$$

where “ \otimes ” is the Kronecker tensor product and (u, v) is the 2D optical flow vector. Then the 2D analysis bank is

$$\text{Low-pass filter: } L = \ell \otimes \ell,$$

$$\text{High-pass filters: } H_1 = \ell \otimes h, H_2 = h \otimes \ell, H_3 = h \otimes h,$$

and the 2D synthesis filter bank is

$$\text{Low-pass filter: } G = g \otimes g,$$

$$\text{High-pass filters: } G_1 = g \otimes q, G_2 = q \otimes g, G_3 = q \otimes q.$$

It is easy to verify that the 2D filter bank defined above is a perfect reconstruction filter bank with the analysis filter bank $\{L, H_i\}$ and the reconstruction filter bank $\{G, Q_i\}$. Then generalizing (12), the iterative equation for the reconstruction of the HR image \tilde{I} from LR images I_k^{LR} amounts to

$$\begin{aligned} \tilde{I}^{n+1} = & \sum_{i=1}^3 Q_i * [(H_i * \tilde{I}^{(n)}) \downarrow_2] \uparrow_2 + G * \frac{1}{K} \left(\sum_{k=1}^K [\tilde{I}_k^{LR} - u \cdot * ((\ell \otimes \tau) * \tilde{I}^n) \downarrow_2 \right. \\ & \left. - v \cdot * ((\tau \otimes \ell) * \tilde{I}^{(n)}) \downarrow_2] \uparrow_2 \right) \end{aligned}$$

Recall that here \tilde{I} is the blurred version of the true I with $\tilde{I} = (c \otimes c) * I$.

There always is noise in the estimated flow u, v . However, the deconvolution operator could make the HR image reconstruction very sensitive to such noise. It is known that the noise variance of the solution will have a hyperbolic growth when the blurring low-pass filter has zeros at high frequency. Thus, de-noising is necessary to suppress the error propagation during the iterative reconstruction.

To suppress the noise, we introduce a wavelet de-noising scheme which subtracts some high-frequency components from \tilde{I}^n . Briefly, we first do a wavelet decomposition of the high-pass response, then apply a shrinkage of wavelet coefficients to the decomposition, and then reassemble the signal.

Our iteration scheme with built-in de-noising operator amounts to

$$\begin{aligned} \tilde{I}^{n+1} = & \sum_{i=1}^3 Q_i * [\Psi(H_i * \tilde{I}^{(n)}) \downarrow_2] \uparrow_2 + G * \frac{1}{K} \left(\sum_{k=1}^K [\tilde{I}_k^{LR} - u \cdot * ((\ell \otimes \tau) * \tilde{I}^n) \downarrow_2 \right. \\ & \left. - v \cdot * ((\tau \otimes \ell) * \tilde{I}^{(n)}) \downarrow_2] \uparrow_2 \right). \end{aligned}$$

The de-noising operator Ψ defined in the equation above is

$$\Psi(H_i * \tilde{I}^n) = G * [(L * (H_i * \tilde{I}^n)) \downarrow_2] \uparrow_2 + \sum_{i=1}^3 [Q_i * (\Gamma[H_i * (H_i * \tilde{I}^n)]) \downarrow_2] \uparrow_2,$$

where Γ is the shrinkage operator.

B. Shrinkage operator and robust regression

The basic idea of wavelet de-noising is to reduce the noise by shrinking the wavelet coefficients where typically most noise exist. Here we take a hybrid shrinkage approach. The hybrid shrinkage operator Γ is defined as:

$$\Gamma(\nu) = \begin{cases} 0 & |\nu| \leq \mu_1; \\ \text{sign}(\nu)\mu_2 \frac{|\nu| - \mu_1}{\mu_2 - \mu_1} & \mu_1 < \nu \leq \mu_2; \\ \nu & \text{Otherwise.} \end{cases} \quad (17)$$

This shrinkage offers benefits both from hard shrinkage (uniformly small risk) and soft shrinkage (overall small risk). The reasoning goes as follows: The optical flow model we adopt here is parametric. The noise introduced by this model is not related to the local intensity variation. A hard shrinkage with appropriate threshold μ_1 should effectively remove such noise. On the other hand, noise in illumination which is related to the intensity of the pixels should be removed adaptively. Therefore, a soft-shrinkage is needed in some range. Finally, in order to keep the sharp edges, we keep the large intensity variations. Occasional outliers will be handled by the median operator when fusing multiple frames.

In summary, our algorithm is as follows: Given an initial HR image $I^{(0)}$ we have

$$I^{(n+1)} = G * \text{median}_k \{ [I_k^{LR} - u_k \cdot * (\ell \otimes \gamma * I^{(n)}) \downarrow_2 - v_k \cdot * (\gamma \otimes \ell * I^{(n)}) \downarrow_2] \uparrow_2 \} + \left(\sum_{i=1}^3 Q_i * [\Phi(H_i * I^{(n)}) \downarrow_2] \uparrow_2 \right). \quad (18)$$

The algorithm above could easily be adapted to different blur filters. We only need to adjust the dual filters G, Q_i to make a new perfect reconstruction filter bank. Also, here we only considered a doubling of the image resolution. But any other resolution increase could be achieved by changing the 2-channel perfect reconstruction filter bank to an M-channel perfect reconstruction filter bank.

C. Relation to regularization methods

One popular de-noising technique used for robust reconstruction is regularization ([19]). Recall that back-projection methods basically find \tilde{x} by minimizing $\sum_{k=1}^K \|y_k - \tilde{y}_k(\tilde{x})\|_2^2$, where $\tilde{y}_k(\tilde{x})$ is the LR signals derived from our estimated \tilde{x} . Such a least squares estimation problem usually is ill-conditioned. One way to increase the stability is to enforce a regularization term

and solve:

$$\min_{\tilde{x}} \sum_{k=1}^K \|y_k - \tilde{y}_k(\tilde{x})\|_2^2 + \alpha \|\Phi(\tilde{x})\|,$$

where Φ is some regularization function and α is some pre-defined smoothing factor. If the regularization is a least squares problem, we call it a Tikhonov-type regularization. The advantage is its simplicity and efficiency, the disadvantage is its relatively poor performance. A nonlinear diffusion regularization, like Total Variation regularization usually performs better, but is computationally expensive.

Wavelet de-noising is closely related to nonlinear diffusion regularization. [15] discussed the relationship of wavelet de-noising to Total Variation regularization for two simple cases. More specifically, consider a wavelet de-noising scheme based on Haar wavelets ($\ell = \frac{1}{2}(1, 1), h = \frac{1}{2}(1, -1)$). Then for the case of the shrinkage operator in the wavelet de-noising being a *soft thresholding* operator defined as in (19), it was shown that the wavelet de-noising process is equivalent to Total Variation based nonlinear diffusion ($\Phi(\tilde{x}) = \|\tilde{x}\|_1$) for a two-pixel signal.

$$\Gamma(\mu) = \begin{cases} \mu - \tau \cdot \text{sgn}(\mu) & \text{if } (|\mu| \geq \tau); \\ 0 & \text{Otherwise} \end{cases} \quad (19)$$

Although [15] only showed the equivalence between Total Variation regularization and a simplistic wavelet de-noising scheme for a signal with 2 pixels, these results still demonstrate that the wavelet de-noising process in our reconstruction is comparable to some nonlinear diffusion regularization schemes in its ability to suppress the error propagation. However, it doesn't have the computational burden of most nonlinear diffusion regularizations, since it only needs a linear wavelet decomposition over one level. In comparison, nonlinear regularizations need to solve a nonlinear optimization.

VI. FLOW ESTIMATION FOR SUPER-RESOLUTION

A. Basic notations

We consider here the planar motion model. In other words, we assume that the underlying 3D structures of the interesting image regions are planar surfaces. Let $I_0, I_1, I_2, \dots, I_K$ be the image frames in the sequence. Fix frame I_0 as the reference image. We need to estimate the homographies between the reference frame I_0 and the frames I_k . There is the following constraint

on the planar homography P_k from reference frame I_0 to frame I_k :

$$P_k = R_k + \vec{v}_k \vec{n}^t, \quad (20)$$

where R_k is a rotation matrix, \vec{v}_k is the translation and \vec{n} is the normal of the plane. For two frames close by, the optical flow \vec{u} at a point \vec{r} is constrained by the brightness consistency constraint ([9])

$$\left(\frac{dI}{d\vec{r}}\right)^t \vec{u}(\vec{r}) = -\frac{dI}{dt}. \quad (21)$$

Let $\vec{p}_k = \text{vec}[P_k] = (p_{1k}, p_{2k}, \dots, p_{9k})^t$ be the vectorized version of the homography P_k . Assume I_k very close to I_0 , then under the small motion assumption, the brightness consistency constraint becomes

$$\left(\frac{dI}{d\vec{r}}\right)^t (\vec{p}_k(\vec{r}) - \vec{r}) = \frac{dI}{dt}, \quad (22)$$

with

$$\vec{p}_k(\vec{r}) = \begin{pmatrix} \frac{p_{1k}x + p_{2k}y + p_{3k}}{p_{7k}x + p_{8k}y + p_{9k}} \\ \frac{p_{4k}x + p_{5k}y + p_{6k}}{p_{7k}x + p_{8k}y + p_{9k}} \end{pmatrix}.$$

$\vec{p}_k(\vec{r})$ is a 2D linear rational polynomial. Thus multiplying the denominator of $\vec{p}_k(\vec{r})$ on both sides of (22) yields a linear homogeneous equation system on \vec{p}_k :

$$A_k \vec{p}_k = 0.$$

B. Multi-frame homography estimation

We need to simultaneously estimate all homographies between the reference image I_0 and the frames I_k . The P_k s are not independent. They share the same plane normal. Using the expression for homographies from [17], we have the following expressions for all homographies which share the same plane normals:

$$P_k = R_k + \vec{v}_k \cdot \vec{n}^t, \quad \text{for } k = 1, \dots, K.$$

In order to improve the estimation of the P_k s, this constraint on the surface normal has to be incorporated into a batch algorithm. Furthermore we need to deal with frames I_k with large displacement to the reference frame I_0 .

Here we take an iterative approach to estimate the homographies P_k . Suppose that at the j th step we are given the approximate solution $P_k^j = R_k^j + v_k^j \cdot (n^j)^t$ for the true homography P_k . Then Equation (22) can be applied to P_k as follows:

$$\left(\frac{dI}{d\vec{r}}\right)^t(\vec{p}_k(\vec{r}) - \vec{p}_k^j(\vec{r})) = I_k(\vec{r}) - I_0(p_k^j(\vec{r})).$$

In other words, we apply the differential brightness consistency constraint between the frame $I_0(p_k^j(\vec{r}))$ (the image obtained by warping I_0 from the homography p_k^j) and $I_k(\vec{r})$. The differential motion is due to the difference between the actual homography and its estimation in the current stage.

We write these linear equations on \vec{p}_k as

$$A_k(P_k^j)\vec{p}_k = 0.$$

Then the minimization across all homographies P_k can be written as

$$\min_{R_k, \vec{n}, v_k} \sum_k \|A_k(P_k^j)vec[R_k + \vec{v}_k \vec{n}^t]\|^2 \quad (23)$$

subject to the constraints that the R_k s are rotation matrices and $\|\vec{n}\| = 1$. This is a constrained minimization, bilinear in R_k , \vec{v}_k and \vec{n} . In the remainder of this section we show how to robustly solve the minimization (23) using an alternative two-steps optimization.

Given P_k^j at step j , we compute $P_k^{j+1} = R_k^{j+1} + \vec{v}_k^{j+1}(\vec{n}^{j+1})^t$ at step $j+1$.

Given $P_k^j = R_k^j + \vec{v}_k^j(\vec{n}^j)^t$, we first update \vec{n}^{j+1} . This minimization of (23) is just a regular least squares minimization over the sphere of \vec{n}^{j+1} and can be written as:

$$\min_{\vec{n}^{j+1}} \sum_k \|A_k(P_k^j)vec[R_k^j + \vec{v}_k^j(\vec{n}^{j+1})^t]\|^2 \quad (24)$$

subject to $\|\vec{n}^{j+1}\| = 1$. The minimization of (24) can easily be solved by SVD decomposition.

The algorithm is as follows: Given A and \vec{b} , the following procedure computes a vector \vec{n} such that $\|A\vec{n} - \vec{b}\|_2$ is minimum, subject to the constraint $\|\vec{n}\| = 1$. Run SVD on A such that $A = U\Sigma V^t$ and save

$$V = \{v_1, v_2, \dots, v_n\}, b = U^t b, \Sigma = diag(\sigma_i).$$

Find λ_* such that

$$\sum_i \left(\frac{\sigma_i b_i}{\sigma_i^2 + \lambda_*}\right)^2 = 1.$$

Then

$$\vec{n} = \sum_i \left(\frac{\sigma_i b_i}{\sigma^2 + \lambda_*} \right) v_i.$$

Then given P_k^j, \vec{n}^{j+1} , we estimate $\{R_k^{j+1}, v_k^{j+1}\}$ for $k = 1, \dots, K$. We need to solve the following minimization: For each k

$$\min_{R_k^{j+1}, v_k^{j+1}} \|A_k(P_k^j) \text{vec}(R_k^{j+1} + \vec{v}_k^{j+1}(\vec{n}^{j+1})^t)\|^2$$

subject to

$$(R_k^{j+1})^t R_k^{j+1} = I_3.$$

This is not a trivial task. Here we present a fast linear method to obtain an approximate solution. Let R_k^{j+1} be approximated by $(I + [\omega]_\times)R_k^j$, where $[\omega]_\times$ is the skew-symmetric matrix of the rotation vector $\vec{\omega}$. Using the observation that

$$P_k^{j+1} = R_k^{j+1} + \vec{v}_k^{j+1}(\vec{n}^{j+1})^t$$

we perform the following decomposition

$$P_k^{j+1} = (I + [\omega_k^j]_\times + \Delta v_k^j \cdot (R_k^j \vec{n}^{j+1})^t) P_k^j,$$

with

$$v_k^{j+1} = (I + [\omega]_x) v_k^j + (1 + (R_k^j \vec{n}^{j+1})^t v_k^j) \Delta v_k^j. \quad (25)$$

Then the minimization (25) is simplified to a standard least squares minimization on ω_k^j and Δv_k^j :

$$\min_{\omega_k^j, \Delta v_k^j} \|A_k(P_k^j) \text{vec}[(I + [\omega]_\times + \Delta v_k^j \cdot (R_k^j \vec{n}^{j+1})^t) P_k^j]\|^2. \quad (26)$$

Thus R_k^{j+1} and v_k^{j+1} can be derived from Δv_k^j and ω_k^j after solving the least squares minimization of (26).

We adopt the procedure in [8] to compute initial values for the homographies P_k^0 s from multiple relative motions P_{k_1, k_2} between frames I_{k_1} and I_{k_2} related by a small displacement. Briefly, the P_k^0 s are determined by an overdetermined linear system:

$$P_{k_1, k_2} P_{k_1}^0 - P_{k_2}^0 = 0.$$

(See [8] for more details.) The iteration algorithm is as follows: Given $P_k^j = R_k^j + \vec{v}_k^j(\vec{n}^j)^t$ at Step j ,

- 1) The minimization (24) is solved by SVD to obtain \vec{n}^{j+1} .
- 2) The minimization (26) is solved by least squares to obtain ω_k^j, Δ_k^j .
- 3) \vec{v}_k^{j+1} is obtained from (25) and the rotation matrix is obtained as:

$$R_k^{j+1} = I + [\omega_k^j]_{\times} + (1 - (1 - \|\omega\|^2)^{\frac{1}{2}})([\omega_k^j]_{\times})^2.$$

The iteration is terminated when the p_k^{j+1} s are close enough to the p_k^j s.

At a quick glance, it seems that decomposing P_k is an overkill since we don't need motion and structure. But actually it doesn't make much difference. It is known that the decomposition of $P = R + \vec{t} \cdot \vec{n}^t$ is unique up to two solutions. By enforcing the consistency between the P_k , the decomposition becomes unique, and it is not difficult to obtain this decomposition.

VII. EXPERIMENTS AND CONCLUSION

We compare our algorithm's high-resolution reconstruction to standard methods using both simulated and real data.

A. Simulated data

We simulated 4 low-resolution images (16×16) from a high resolution image by shifting, blurring and downsampling. The blurring filter is

$$\ell = \frac{1}{16} \begin{pmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{pmatrix}.$$

Three kinds of noise were simulated:

- 1) Error in motion estimation. This error is modeled by local Gaussian white noise with parameter σ . The local covariance matrix is due to the magnitudes of the image gradients.
- 2) Noise in pixel formation. We added Gaussian white noise with parameter γ to the pixel values.
- 3) Error in PSF modeling. We also checked how error in the PSF modeling influences the performance. The approximated PSF $\hat{\ell}$ used in the reconstruction was

$$\hat{\ell} = \frac{1}{16} \begin{pmatrix} 1 & 1 & 1 \\ 1 & 8 & 1 \\ 1 & 1 & 1 \end{pmatrix}. \quad (27)$$

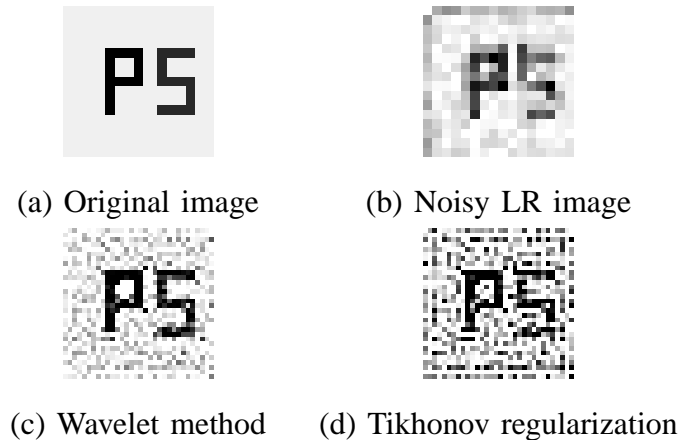


Fig. 3. The HR images (c) and (d) are reconstructed from four LR images by five iterations. (c) is reconstructed by our method. (d) is reconstructed by the back-projection method with Tikhonov regularization. The motion noise is local Gaussian noise with $\sigma = 0.2$. The image formation noise is Gaussian noise with $\gamma = 0.01$. The approximation $\hat{\ell}$ in (27) is used in the reconstruction instead of the true PSF ℓ .

We compared our wavelet-based method to the popular “POCS” back-projection method ([20]) enforced by Tikhonov regularization (See Fig. 3). It may be possible that another scheme, namely Total Variation regularization would give a bit better results. However, this would require solving a nonlinear minimization over each iterative step during the reconstruction, which is computationally expensive. In our implementation the regularization term is the 2-norm of the Laplacian smoothness constraint with roughly tuned regularization parameter α


Fig. 4 demonstrates the performance of the wavelet-based method for various noise settings. Performance is measured by the SNR (Signal-to-Noise ratio) of the reconstructed image to the true image, which is defined as:

$$SNR = 20 \log_{10} \frac{\|x\|_2}{\|x - \hat{x}\|_2},$$


where \hat{x} is the estimate for the true image x . Fig. 4 clearly indicates the advantage of our wavelet-based method in suppressing the noise. Especially when noise is large, the boost in performance is significant.

B. Real data


We used an indoor sequence depicting 13 image frames of a paper box (Fig. 5). The sequence was taken with a hand-held camera. An interesting planar region was chosen manually. Fig. 6



(a) Comparison for motion noise



(b) Comparison for image formation noise



(c) Comparison for PSF error

Fig. 4. Comparison between the our method and POCS with Tikhonov regularization for various amounts of noise. The reconstructed image has been obtained in 5 iterations. The x-axis denotes the variance of the noise, the y-axis denotes the SNR of the reconstruction.

and Fig. 7 show a comparison of the results from four different methods for different regions. Here the reconstructed HR images double the resolution of the LR images. The HR image in Fig. 6-7(a) were obtained by cubic interpolation from a single LR image. In Fig. 6-7(b) we used the POCS method, and estimated the flow field using an affine motion model. In Fig. 6-7(d) we used the POCS method, and estimated the flow field with the Homography model. Fig. 6-7(d) show the results from our reconstruction scheme. The difference can be visually evaluated. As can be seen, there is improvement from (b) to (c) and from (c) to (d) in Fig. 6 and Fig. 7. The letters in Fig. 6(d) and Fig. 7(d) are the clearest, and there are minimal artifacts around the edges.




(a) Reference frame (b) LR planar image region


Fig. 5. Reference image frame of first indoor video and its selected region.



(a) Interpolation



(b) POCS+affine



(c) POCS+Homography




(d) Wavelet+Homography


Fig. 6. Comparison of one reconstructed HR region for various methods.




(a) Interpolation



(b) POCS + Affine



(c) POCS + Homography



(d) Wavelet + Homography

Fig. 7. Comparison of another reconstructed HR region from Fig. 5 for various methods.



Fig. 8. Reference frame from second indoor video.

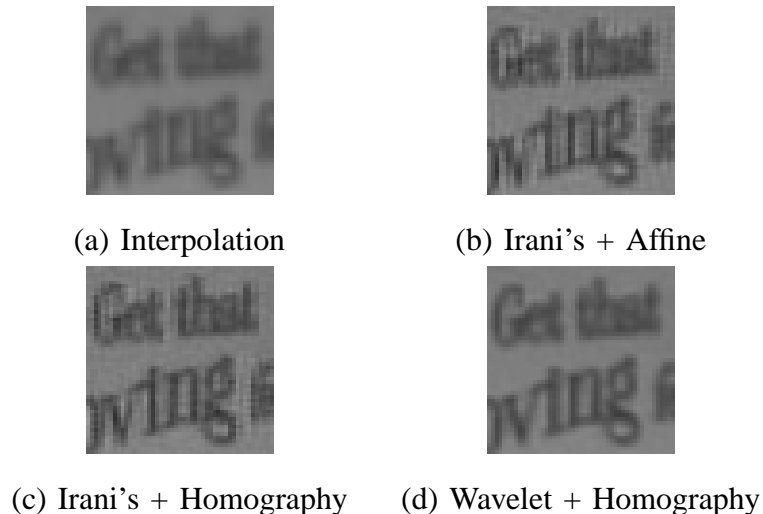


Fig. 9. Comparison of one reconstructed HR region from Fig. 8 for various methods.

A second indoor sequence depicting a box wrapped in newspaper (Fig. 8) was tested. This time, we compared our method against Irani's method [10]. See Fig. 9 for a visual comparison. The same conclusion holds as for the previous experiment; both our homography-based motion estimation and our reconstruction method lead to improved results.

We also used an outdoor sequence of 11 frames containing a warning sign in the scene. We compared the reconstructions from different methods for a manually selected region. Since the regions are too small to provide enough information for the homography flow model, here instead we used an affine flow model. See Fig. 10 for a visual comparison.

In recent years an effort started combining geometric constraints with Signal Processing [7], [12], [13]. Along these lines we have presented an algorithm for the problem of super-resolution reconstruction. We presented a new method for estimating the homography between multiple frames in a sequence, and a new wavelet-based reconstruction algorithm. We demonstrated both




Fig. 10. (a) The key frame in the video. (b) The reconstruction from the interpolation. (c) The reconstruction from Irani's method using affine flow. (d) The reconstructed image from the wavelet method with de-noising using affine flow.

in theory and experiments that the proposed method is very robust to noise without sacrificing efficiency. The reconstruction scheme allows for super-resolution reconstruction from general video sequences, even when the estimated optical flow is very noisy, and it outperforms existing methods.

REFERENCES

- [1] S. Baker and T. Kanade. Limits on super-resolution and how to break them. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 372–379, 2000.
- [2] B. Bascle, A. Blake, and A. Zisserman. Motion deblurring and super-resolution from an image sequence. In *Proc. European Conf. Computer Vision*, volume 2, pages 573–582, 1996.
- [3] A. Chambolle, R. Devore, N. Lee, and B. Lucier. Nonlinear wavelet image processing: Variational problems, compression and noise removal through wavelets. *IEEE Trans. Image Processing*, 7, 1998.
- [4] R. Chan, T. Chan, L. Shen, and Z. Shen. Wavelet deblurring algorithms for spatially varying blur from high-resolution image reconstruction. *Linear Algebra and its Applications*, 366:139–155, 2003.
- [5] M. Eland and A. Feuer. Restoration of a signal super-resolution image from several blurred, noisy and undersampled measured images. *IEEE Trans. on Image Processing*, pages 1646–1658, 1997.
- [6] S. Farsiu, D. Robinson, M. Elad, and P. Milanfar. Robust shift and add approach to super-resolution. In *SPIE*, 2003.
- [7] C. Fermüller, P. Baker, and Y. Aloimonos. Visual space time geometry: A tool for perception and imagination. *Proceedings of the IEEE*, 90(5):1113–1135, 2002.
- [8] V. M. Govindu. Combining two-view constraints for motion estimation. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2001.
- [9] B. K. P. Horn. *Robot Vision*. MacGraw Hill, New York, 1986.

- [10] M. Irani and F. Peleg. Motion analysis for image enhancement: Resolution, occlusion and transparency. *Journal of Visual Comm. and Image Repr.*, 4:324–335, 1993.
- [11] H. Ji and C. Fermüller. A 3d shape constraint on video. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 28(6):1018–1023, 2006.
- [12] A. Makadia and K. Daniilidis. Rotation estimation from spherical images. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 28:1170–1175, 2006.
- [13] A. Makadia, C. Geyer, and K. Daniilidis. Correspondenceless structure from motion. *International Journal of Computer Vision*, 2007.
- [14] S. Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, 1999.
- [15] P. Mrazek, J. Weickert, and G. Steidl. Correspondences between wavelet shrinkage and nonlinear diffusion. In *Proc. Scale-Space Methods in Computer Vision*, pages 101–116, 2003.
- [16] N. Nguyen and N. P. Milanfar. A wavelet-based interpolation-restoration method for superresolution. *Circuits, Systems and Signal Processing*, 19:321–338, 2000.
- [17] A. Shashua and S. Avidan. The rank 4 constraint in multiple (≥ 3) view geometry. In *Proc. European Conf. Computer Vision*, 1996.
- [18] A. Tekalp, M. Ozkan, and M. Sezan. High-resolution image reconstruction from low-resolution image sequences and space-varying image restoration. In *ICASSP*, pages 169–172, 1992.
- [19] J. Weickert. *Anisotropic Diffusion in Image Processing*. Teubner, Stuttgart, 1998.
- [20] C. Youla. Generalized image restoration by the method of alternating orthogonal projections. *IEEE Trans. Circuits System*, 1978.
- [21] L. Zelnik-Manor and M. Irani. Multi-frame estimation of planar motion. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 22:1105–1114, 2000.
- [22] W. Zhao and H. S. Sawhney. Is super-resolution with optical flow feasible? In *Proc. European Conf. Computer Vision*, pages 599–613, 2002.
- [23] A. Zomet, A. Rav-Acha, and S. Peleg. Robust super-resolution. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2001.