The Translation Sensitivity of Wavelet-Based Registration

Harold S. Stone, *Fellow*, *IEEE*,
Jacqueline Le Moigne, *Senior Member*, *IEEE*,
and Morgan McGuire, *Member*, *IEEE*

Abstract—This paper studies the effects of image translation on wavelet-based image registration. The main result is that the normalized correlation coefficients of low-pass Haar and Daubechies wavelet subbands are essentially insensitive to translations for features larger than twice the wavelet blocksize. The third-level low-pass subbands produce a correlation peak that varies with translation from 0.7 and 1.0 with an average in excess of 0.9. Translation sensitivity is limited to the high-pass subband and even this subband is potentially useful. The correlation peak for high-pass subbands derived from first and second-level low-pass subbands ranges from about 0.0 to 1.0 with an average of about 0.5 for Daubechies and 0.7 for Haar. We use a mathematical model to develop these results, and confirm them on real data.

Index Terms—Image registration, image search, wavelet transform, correlation, low-pass filter, high-pass filter, Haar wavelet, Daubechies wavelet, multiresolution search.

1 Introduction

IMAGE registration is a core tool for studying how images differ. By aligning the features of one image with the features of another, pixel-by-pixel differences reveal much information about what has changed between the two images. This simple operation, as well as more sophisticated operations that build upon it, is the key tool in the study of earth's environment from satellite images and medical pathologies from radiological images. The growing volume of images to study and the cost of registration place a high potential value on efficient and robust registration algorithms. For surveys of image-registration research, see [1], [2]. This paper examines a high-speed image registration technique in detail to determine its sensitivity to translations of images.

Many different registration techniques have been proposed. They fall mainly into three categories:

1. feature identification, followed by feature matching according to models that best characterize how one image differs from another [3], [4], [5],
2. phase coherence in the Fourier domain to determine the most likely rigid translation, rotation, and scale change [6], [7], [8],
3. full or low-resolution pixel correlations to determine the most likely rigid translation, rotation, or scale change [9], [10], [11], [12], [13], [14].

Registration can be very costly depending on the size and resolution of images and the type of mapping that carries one image into another. In this paper, the mappings considered consist of the set of rigid translations in the X-Y plane. An image of size N pixels has O(N) possible translations to consider, each of which requires O(N) operations to compute a correlation with a candidate image, thus producing an overall complexity of O(N^2) when performed in the pixel domain. A reduction in resolution by a factor of k reduces the cost of the image registration by a factor of k^2. (Fourier methods for computing low-resolution correlations may be even faster. [15].) Several wavelet-based registration methods have been developed where features are high-interest points or local curvature discontinuities from a wavelet representation [2], [13], [16], [17], [18], [19], [20], [21]. Our study is based on multilevel correlation of selected wavelet features [2], [13], [16], [20]. These papers give various examples of searches performed on low-resolution images, which are subsequently refined to high-resolution precision by local search. Allen et al. use a cubic spline wavelet and a sum of squares criterion, whereas the other three references use Daubechies wavelets and normalized correlation in their studies. Allen et al. show how to deal with multiple candidate matches through the refinement process.

It is well-known that wavelet representations are not translation invariant [16], [22], [23]. Consequently, wavelet-based registration at low and high resolution may produce inconsistent results. This observation has led to the development of translation-invariant wavelets [22], [23], [24], [25], [26], which necessarily lead to higher computational complexity in registration algorithms and, thus, are at odds with the goal of gaining faster registration by reducing the number of points to register.

This paper examines the sensitivity of wavelet-subband correlations to translation. We use the Haar wavelet and the four-tap Daubechies wavelet [27]. The two outputs of a one-dimensional (1D) wavelet filter are the *low-pass* and *high-pass* subbands. The original image is a Level-0 wavelet. Level 1 of the wavelet contains the low-pass and high-pass subbands of the original image. The Level-d subbands are the low-pass and high-pass subbands of the Level-$d-1$ low-pass subband. The *blocksize*, B, of a Level-d subband is the input offset (in pixels) for successive outputs in that subband. Blocksize $B = 2^d$ in the two-subband model. Note that B is equal to the resolution reduction factor of a wavelet filter. For two-dimensions (2D), a wavelet has four subbands at each level. They are the four combinations of low and high-pass filtering of rows and columns. The low-pass Level-d subband is the low-low subband of the low-low subband at Level $d-1$. The high-pass subband used in this paper is the low-high subband of the low-low subband at Level $d-1$. Blocksize $B = 2^{2d}$ for 2D Level-d subbands.

This paper shows that the low-pass subband is essentially insensitive to translation for feature sizes at least twice the block size B of the subband. Note that this is consistent with the Sampling Theorem because if a low-pass subband contains at least two samples of a feature in all translations, that feature is stable with respect to translation in registration correlations.

A mathematical model shows that the low-pass correlation peak is as low as 0.8 in 1D and 0.7 in 2D for features twice the blocksize or more. The average height of the low-pass correlation peak exceeds 0.9 for both 1D and 2D models. The high-pass correlation peak is sensitive to translation, however, and can fall below 0.3 in 1D and below 0.1 in 2D for some translation phases. Nevertheless, the average correlation peak is 0.5 or higher for both 1D and 2D, which indicates that high-pass correlation can be used, but with care. For example, coarse-to-fine registration can carry forward multiple candidates [16].

To validate the mathematical model, the paper presents the results of experiments conducted on a real image and shows that the translation effects observed in practice are consistent with those exhibited by the mathematical model. We also derive a lower bound based on white-noise statistics. This model supports the observation that the low-pass subband is insensitive to translation.

- *H.S. Stone is with the NEC Research Institute, 4 Independence Way, Princeton, NJ 08540. E-mail: hstone@research.nj.nec.com.*
- *J. Le Moigne is with the Applied Information Science/CS Branch, Code 935, NASA Goddard Space Flight Center, Greenbelt, MD 20771. E-mail: lemoigne@backserv.gsfc.nasa.gov.*
- *M. McGuire is with the Massachusetts Institute of Technology, Cambridge, MA 02139. E-mail: morgan3d@yahoo.com.*

Manuscript received 12 Feb. 1998; revised 29 June 1999.
Recommended for acceptance by R. Chellappa.
For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number 107621.

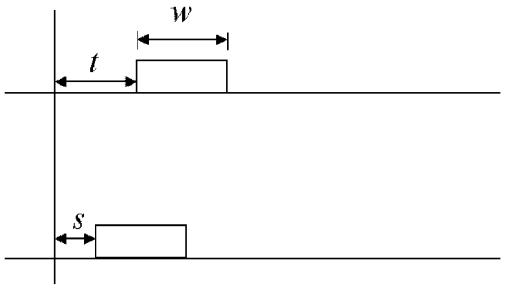

Fig. 1. The model function used for registration studies.

Section 2 describes the mathematical model and white-noise lower bound. Section 3 compares the mathematical results with the results obtained for a test image. The last section contains a summary and suggestions for future research.

2 Modeling of Translation Effects

This section contains an idealized mathematical model that enables us to study translation sensitivity. A white-noise model, which serves as a worst-case bound, follows the idealized model.

2.1 Pulse-Model Experiments

The 1D signal model appears in Fig. 1. It is a pulse of unit height and width w offset from the origin by an amount t. It is correlated against a pulse of identical width offset by an amount s. In the absence of noise, the correlation peak occurs at a point offset from the origin by an amount whose absolute value is $|s-t|$. The model uses a baseline length of 128 and varies pulse widths from 1 to 64. The study covers all distinct different values of $|s-t|\ mod\ B$, where B is the filter blocksize. For two dimensions, the pulse model becomes a square mesa. A phase shift of a mesa by amount s moves it by s in both the x and y coordinates. This reduces the number of phase shifts to study from B^2 to B.

Fig. 2 illustrates the reason for translation effects. A sample signal appears twice in the upper half of the figure. The right-hand copy is shifted by two relative to the left-hand copy. The encircled points and the arrows illustrate the result of filtering the signal with a filter whose block size is 4, leaving a filtered result shown in the lower half of the figure. The effect of translation causes different sets of four points to be processed by the filter, thereby producing slightly different filtered versions of the original signal. When we try to register the filtered versions of the signals to each other, the discrepancy in shape produces a potential registration error.

Our registration criterion for this study is the normalized correlation coefficient defined as follows: Let $\mathbf{x}$ and $\mathbf{y}$ be column N-vectors with indices in the range $0 \le i \le N-1$. The function $C(\mathbf{x},\mathbf{y})$, the *normalized correlation coefficient* of x and y, is:

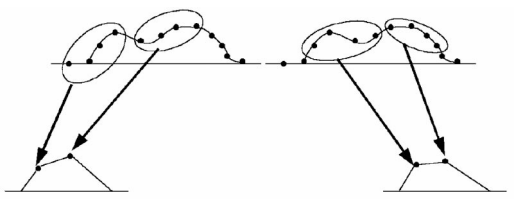

Fig. 2. Translation effects on a wavelet filtered version of a signal.

$$C(\mathbf{x},\mathbf{y}) = \frac{\frac{1}{N}\sum_{i=0}^{N-1} x_i y_i - \left(\frac{1}{N}\sum_{i=0}^{N-1} x_i\right)\left(\frac{1}{N}\sum_{i=0}^{N-1} y_i\right)}{\sqrt{\left(\frac{1}{N}\sum_{i=0}^{N-1} x_i^2 - \left(\frac{1}{N}\sum_{i=0}^{N-1} x_i\right)^2\right)\left(\frac{1}{N}\sum_{i=0}^{N-1} y_i^2 - \left(1N\sum_{i=0}^{N-1} y_i\right)^2\right)}}. \quad (1)$$

The registration problem in one dimension is to find the translation of vector y that maximizes the correlation coefficient. We define $\mathbf{y}^{(j)}$ to be the N vector created by translating $\mathbf{y}$ j places to the right, and while shifting on 0s from the left. Given model functions $\mathbf{x}$ and $\mathbf{y}$, containing a pulse of width w starting at positions s and t, respectively, at full resolution the best registration position is the value of j that maximizes $C(\mathbf{x},\mathbf{y}^{(j)})$ for $-N \le j \le N$. The value of the peak correlation is 1 and $|j| = |s-t|$.

Let $\mathbf{F}$ denote a wavelet filter such that $\mathbf{Fx}$ is a wavelet representation of $\mathbf{x}$ reduced in resolution by a factor of B, the filter blocksize. To use wavelet registration successfully, we want $C(\mathbf{Fx},\mathbf{F}(\mathbf{y}^{(j)}))$ to take on its maximum value at a value of $|j|$ close to $|s-t|/B$. Since we register by choosing the value of j that maximizes $C(\mathbf{Fx},\mathbf{F}(\mathbf{y}^{(j)}))$, we want to assure that incorrect registrations produce low correlations and the correct registration position produces a correlation close to unity.

Figs. 3 and 4 show the statistics on the height of the correlation peak as a function of pulse width for the low-pass and high-pass filters, respectively. The top row of plots in each figure show the minimum height of the peak over all pairs of phases s and t. The bottom row shows the average height of the correlation peak. The left and right columns plot 1D and 2D data, respectively.

Recall that the 2D model treats only those phase shifts for which a pulse moves by the same amount in both directions. These phase shifts include both the best and worst cases. The average peak height over these phase shifts tends to be a little higher than the average taken over all distinct phase shifts.

The high-pass filter data in Fig. 4 had to be stabilized for the Haar filter because, for some pulse widths and offsets, the Haar high-pass output contains no energy. The correlation coefficient is undefined for this case because both the numerator and denominator of (1) are 0. Of the various ways to stabilize the computation, we chose to stabilize by adding a small amount of noise to the idealized signal. The plots in Fig. 4 are for the correlations of a noisy pulse with a shifted version of itself. The additive noise has a standard deviation of 0.05 and the pulses are of height 1.00, so the noise is a very small perturbation of the signal. In a separate experiment, we stabilized the correlation by lengthening the pulse edges to rise and fall over two time intervals instead of within a single interval. Both stabilization methods gave similar results in that they did not perceptibly alter the stable correlations and they produced similar answers for points where the computation was formerly unstable. We also verified that similar noise does not perturb the low-pass correlations perceptibly.

The important information in Figs. 3 and 4 is that the low-pass filter produces very high correlation peaks in the worst-case and on the average for features greater than $2B$ in size. The high-pass filter produces fairly good correlation peaks on average, but very poor in the worst case. Experiments carried out at Levels 1 and 2 not illustrated here showed the same characteristics. The worst-case correlations in Fig. 4 confirm the findings in the literature about the translation sensitivity of wavelets. But, the average data indicates that the high-pass subband contains useful information that can aid registration in many cases, if not in all cases. A correlation 0.6 in the high-pass subband is a good indicator of a correct registration. The main problem is that, in the worst-case, the correlation peak might be very low and might be lost among taller false peaks.

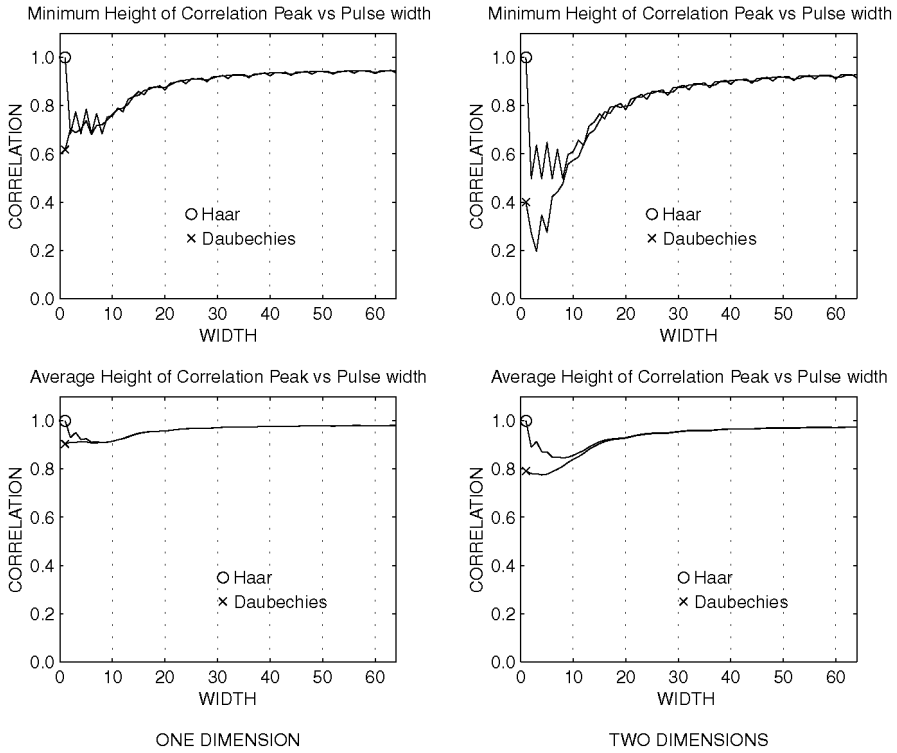

Fig. 3. Minimum and average peak correlation heights as a function of pulse width. These are for Level-3 low-pass subbands.

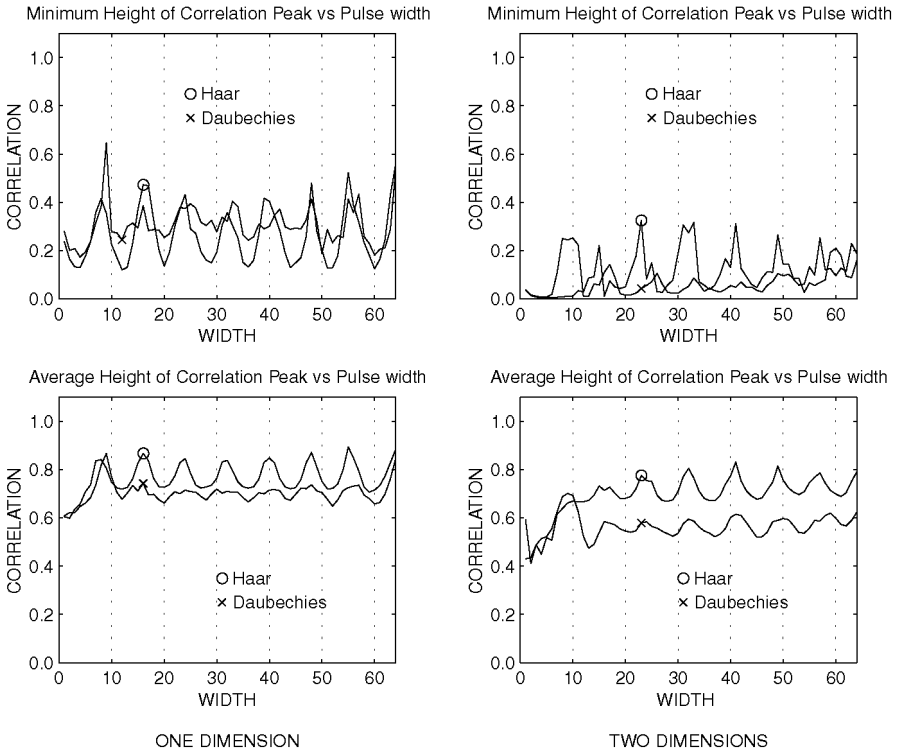

Fig. 4. Minimum and average peak correlation heights as a function of pulse width. These Level-3 subband correlations are for the high-pass subbands of Level-2 low-pass subbands. The Haar data have been stabilized by adding small amounts of noise to the pulses.

The periodic structure of the correlation peak heights in Fig. 4 is due to the block structure of the wavelets. For a fixed offset s and varying width w, the wavelet coefficients of the right-hand edge of the pulse vary with a period of B, the block size of the wavelet subband.

2.2 White-Noise Analysis

The pulse-model results are partially due to the way the filter coefficients correlate with each other and partially due to image features that the filters amplify. This section uses white noise to reveal how filter coefficients affect correlations. Normalized correlations of shifted versions of white noise serve as a lower

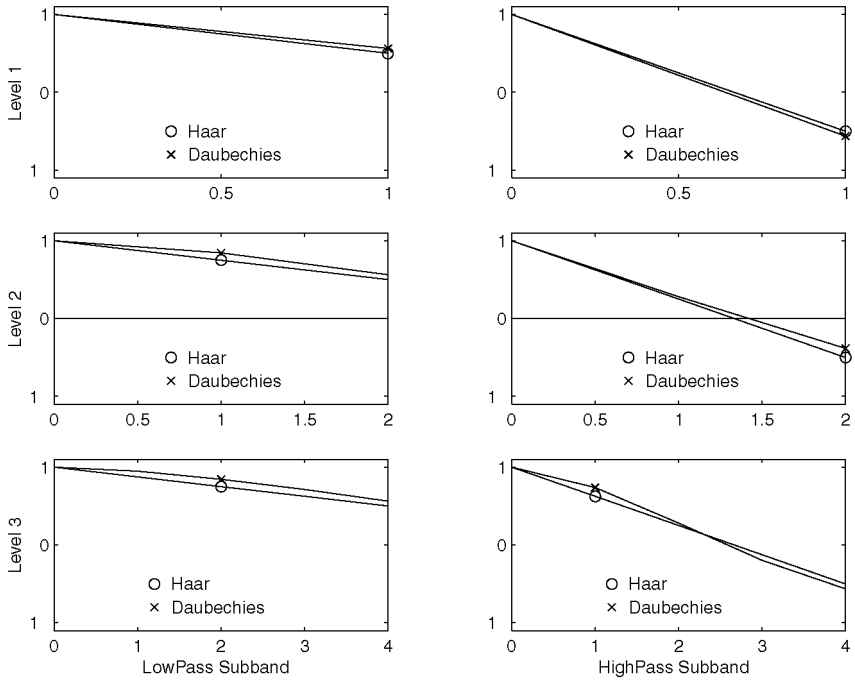

Fig. 5. Truncated correlations of Daubechies and Haar filter coefficients.

bound on the correlation function of shifted versions of real images. Real images have strong correlations among neighboring pixels. White noise, by definition, has no correlations among neighbors.

For this analysis, we let our "image" be a vector of N iid Gaussian variates with mean 0 and variance σ^2. Let $\mathbf{F}$ be a wavelet-subband filter with block size B and $\mathbf{Fx}$ be the filtered subband of a white-noise N-vector $\mathbf{x}$. The length of $\mathbf{Fx}$ is $M = N/B$. To deal with the special case of the last wavelet coefficient, which for the Daubechies filter depends on components of $\mathbf{x}$ with indices greater than N, we assume that the filter $\mathbf{F}$ is circularly symmetric. Hence, subscripts of $\mathbf{x}$ wrap around modulo N. A different model is to treat the noise beyond the end of the image as being identically 0, which is an end-off data shift instead of cyclical shift. Later, we show cyclical shifts produce essentially the same results as end-off shifts for the vector lengths and filters of interest.

We can rewrite (1) in the form:

$$C\Big(\mathbf{Fx}, \mathbf{F}\Big(\mathbf{x}^{(j)}\Big)\Big) = \frac{(\mathbf{Fx})^T[\mathbf{F}\big(x^{(j)}\big) - \big(\frac{1}{M}\big)\big(\sum \mathbf{Fx}\big)\big(\sum \mathbf{F}\big(\mathbf{x}^{(j)}\big)\big)}{\sqrt{\Big((\mathbf{F}(\mathbf{x}))^T\mathbf{F}(\mathbf{x}) - \big(\frac{1}{M}\big)\big(\sum \mathbf{Fx}\big)^2\Big)\Big((\mathbf{F}(\mathbf{x}^{(j)}))^T\mathbf{F}(x^{(j)}) - \big(\frac{1}{M}\big)\big(\sum \mathbf{F}(\mathbf{x}^{(j)})\big)^2\Big)}}, \quad (2)$$

where $\sum \mathbf{Fx}$ denotes the sum of the components of M-vector $\mathbf{Fx}$. We assume that $E\big[C\big(\mathbf{Fx}, \mathbf{F}\big(\mathbf{x}^{(j)}\big)\big)\big]$ is approximately equal to the ratio of the expected values of the numerator and denominator of (2). (Experiments confirmed this approximation for Gaussian noise.)

Recall that the block size B of a filter is the resolution-reduction factor. Let L be the number of taps for a filter. For Haar filters, $L = B$. Daubechies 4-tap filters at Level 2 and 3 have $L = 10$ and $B = 4$, $L = 22$ and $B = 8$, respectively. The mean of the first numerator term is:

$$\begin{aligned} &E\Big[(\mathbf{Fx})^T\mathbf{F}\Big(\mathbf{x}^{(j)}\Big)\Big] = \\ &E\left[\sum_{i=0}^{M-1}\left(\sum_{k=0}^{L-1} f_k x_{iB+k} \sum_{r=0}^{L-1} f_r x_{iB+r+j}\right)\right] = \\ &E\left[\sum_{i=0}^{M-1}\left(\sum_{k=0}^{L-1} f_k x_{iB+k} \sum_{r=j}^{L-1+j} f_{r-j} x_{iB+r}\right)\right], \end{aligned} \quad (3)$$

where all subscripts of $\mathbf{x}$ are calculated modulo N. Note how the block size B causes the filter to act on components of $\mathbf{x}$ at indices B apart. Because $E\big[x_i x_j\big] = 0$ for $i \neq j$ and is σ^2 otherwise, we can simplify (3) to the form:

$$\begin{aligned} E\Big[(\mathbf{Fx})^T\mathbf{F}\Big(x^{(j)}\Big)\Big] &= E\left[\sum_{i=0}^{M-1}\left(\sum_{k=j}^{L-1} f_k f_{k-j} x_{iB+k}^2\right)\right] \\ &= M\sigma^2 \sum_{k=j}^{L-1} f_k f_{k-j}. \end{aligned} \quad (4)$$

The summation is a truncated correlation of the filter coefficients. Fig. 5 plots these for the range of offsets from 0 to $B/2$, the range of interest.

The second term in the numerator involves factors of the form $\sum \mathbf{Fx}$, which expands into

$$\begin{aligned} &E\Big[\Big(\sum \mathbf{Fx}\Big)\Big(\sum \mathbf{F}\Big(\mathbf{x}^{(j)}\Big)\Big)\Big] = \\ &E\left[\left(\sum_{i=0}^{M-1}\sum_{k=0}^{L-1} f_k x_{iB+k}\right)\left(\sum_{r=0}^{M-1}\sum_{s=0}^{L-1} f_s x_{rB+s+j}\right)\right]. \end{aligned} \quad (5)$$

Because the coefficient of each x_i in (5) is a sum of filter coefficients, we define the B-tap $\mathbf{G}$ filter to be $g_i = \sum_{j=0}^{\lceil L/B \rceil} f_{i+jB}$. Replacing sums of $\mathbf{F}$ coefficients in (5) with $\mathbf{G}$ coefficients yields

$$\begin{aligned} &E\Big[\Big(\sum \mathbf{Fx}\Big)\Big(\sum \mathbf{F}\Big(\mathbf{x}^{(j)}\Big)\Big)\Big] = \\ &E\left[\left(\sum_{i=0}^{N-1} x_i g_{i \bmod B}\right)\left(\sum_{k=0}^{N-1} x_k g_{(k-j) \bmod B}\right)\right]. \end{aligned} \quad (6)$$

Since $E[x_i x_k] = 0$ for $i \neq k$, (6) reduces to

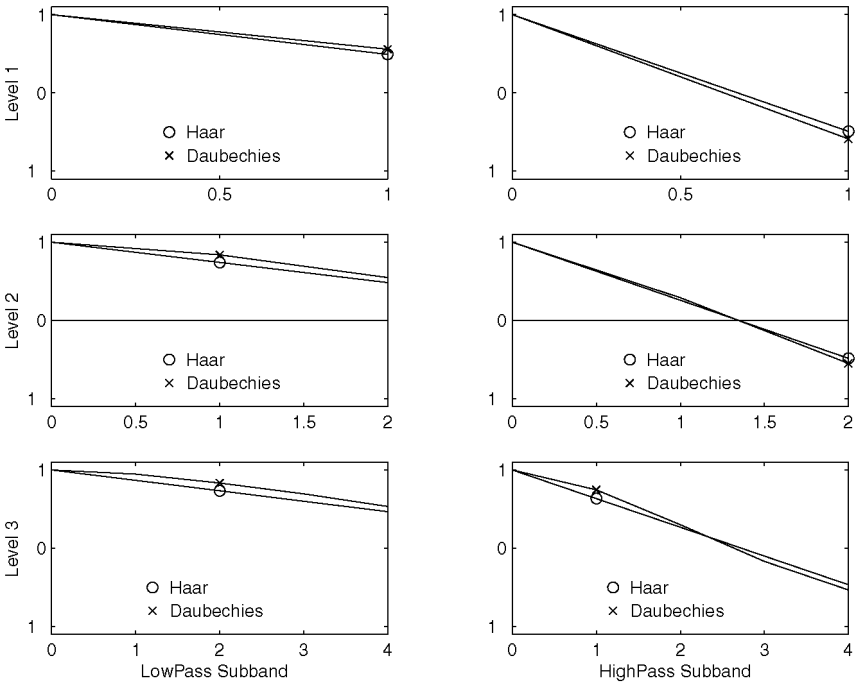

Fig. 6. Normalized correlations of wavelets of white noise.

$$E\Big[\Big(\sum \mathbf{Fx}\Big)\Big(\sum \mathbf{F}\Big(\mathbf{x}^{(j)}\Big)\Big)\Big] = E\left[\sum_{i=0}^{N-1} x_i^2 g_{i \bmod B}\, g_{(i-j) \bmod B}\right] = M\sigma^2 \sum_{i=0}^{B-1} g_{i \bmod B}\, g_{(i-j) \bmod B}. \quad (7)$$

The last line of (7) is a circular correlation of $\mathbf{G}$ which repeats M times in a vector of length $N = BM$. The $\mathbf{G}$ filter coefficients are constant and equal to $1/\sqrt{B}$ for both the Haar and Daubechies low-pass filters, so the circular cross-correlation of these filters is equal to unity in that subband.

The denominator of (2) has two factors similar to those in the numerator. Recall that the first factor has a mean of $(M-1)\sigma^2$. The second factor, $\mathbf{F}\big(\mathbf{x}^{(j)}\big)^T \mathbf{F}\big(\mathbf{x}^{(j)}\big) - (1/M)\big(\sum \mathbf{Fx}^{(j)}\big)^2$, has an identical mean because it is a cyclic shift of the first factor.

Fig. 6 plots an estimate of the correlations obtained by replacing each of the terms in (2) by their means. Note how closely Fig. 6 resembles the noise-free wavelet coefficient correlations in Fig. 5. A close approximation of the low-pass correlation for phase shifts up to $B/2$ is $(B-s)/B$, where s is the absolute difference in phase. The ratio $(B-s)/B$ is the fraction of overlapping components for phase shift s. The worst-case correlation occurs at an offset of size $|B/2|$ and is approximately 0.5 for all low-pass filters in Fig. 6. For 2D, a good approximation to the low-pass correlation is $(B-s)(B-t)/B^2$, where s and t are the absolute phase differences in each dimension. The lowest correlation occurs when $s = t = B/2$ and is approximately 0.25.

The assumption that we can use cyclic shifts to approximate end-off shifts is a good one for white noise. Consider the terms affected by this assumption in (1). The cross-correlation in the numerator is unchanged as long as no noise component shifts off and wraps around into the same filter block. This does not occur for small shifts of long vectors, and occurs nowhere in data plotted. The factor $\mathbf{Fx}^{(j)}$ in the second numerator term has a diminishing mean for end-off shifts and a constant mean for cyclical shifts. This difference is very small when the size of the shift is small compared to the length of the vector. (We shifted up to 4 positions of 128 in a vector.) Since the first term is O($M\sigma^2$) and the second is O(σ^2), small changes in the second term have very little effect overall. The decrease is balanced in part by a decrease in the variance term for $\mathbf{Fx}^{(j)}$ in the denominator.

2.3 Understanding the Behavior of Wavelet Correlation

The main result is that low-pass subbands are relatively insensitive to translations. The white-noise model indicates that translation alone cannot lower the peak below 0.25 for real images. The pulse model predicts even lower effects for real images.

To understand why translation sensitivity is so low, consider a wavelet coefficient $f_0x_i + f_1x_{i+1} + \ldots + f_{L-1}x_{i+L-1}$ and the closest corresponding wavelet coefficient of the signal offset by j positions, which is $f_0x_{i+j} + f_1x_{i+1+j} + \ldots + f_{L-1}x_{i+L-1+j}$. For small j, $\mathbf{Fx}$ and $\mathbf{F}\big(\mathbf{x}^{(j)}\big)$ have identical wavelet coefficients except for the wavelet coefficients near the boundary of the pulse. Therefore, most of the terms in (2) are the same for both in-phase and shifted correlations. The exceptions are the terms that span the edges of the pulse. If a pulse is sufficiently large to be visible in the low-pass subband of every cyclic shift of the pulse, the exceptional wavelet coefficients degrade the normalized correlation by relatively small amounts. A narrow pulse may not be visible in some subbands, and its energy contributes only to the exceptional terms. This leads to nonnegligible degradation of the correlation coefficient. If a pulse has width $2B-2$ or less, there is at least one cyclic shift of the pulse for which no wavelet coefficient depends on at least B samples from the pulse. A pulse of width $2B-1$ contains at least B samples that contribute to a single wavelet coefficient in every cyclic shift. The data in Fig. 3 clearly show how the correlation climbs close to unity as the pulse width increases above $2B$.

The high-pass subband behaves somewhat differently. The Haar and Daubechies high-pass filter coefficients have a large derivative near the middle tap and highlight any edge that spans this region. An edge that occurs anywhere else within the block is attenuated relative to a highlighted edge. (The Haar high-pass filter totally eliminates edges that occur between blocks.) Nevertheless, the pulse model suggests that high-pass coefficients work moderately well on average.

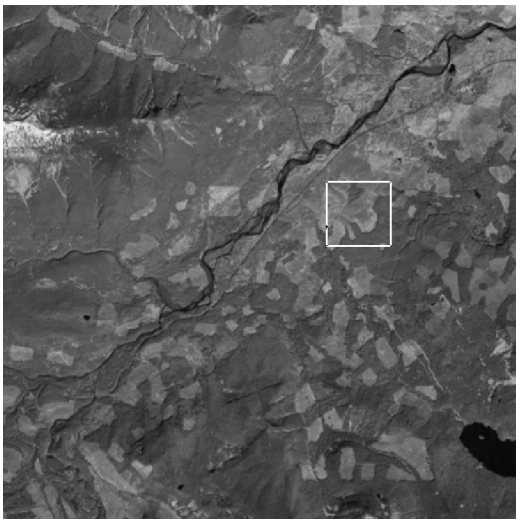

Fig. 7. Satellite image and template from the image to be registered.

3 Experiments with a Real Image

To validate the observations of the previous section, in this section we show that the translation sensitivity of wavelet-subband correlations of a real image are consistent with our mathematical models. Fig. 7 shows the satellite image of a section of the state of Washington, acquired by the Landsat Thematic Mapper (TM). The original image is 512×512 and the highlighted search template is 64×64. Our experiment uses Level-3 wavelets, which creates an image of size 64×64 and a template of size 8×8. For both Haar and Daubechies filters, we created different wavelets for each of the 64 possible phases of the template relative to an 8×8 grid. Because the image also has 64 different phases relative to the same grid, there are 4,096 distinct relative phases to treat. Although we examined 64 of them, our data includes the best and worst-case relative shifts.

The previous section suggests that features to be registered should be at least $2B$ in size, which is 16×16 pixels for $B = 8$. The light large peninsular feature in the center of the template is approximately 20×20 pixels. The small peninsular region to its left may be invisible in the low-pass subband for some offsets.

The experimental results for the low-pass subband appear in Figs. 8 and 9. These three-dimensional plots show the peak correlation as a function of the x and y offsets of the template with respect to the wavelet grid. The maximum correlation value is unity in these figures. The Haar wavelet in Fig. 8 and the Daubechies wavelet in Fig. 9 exhibit essentially the same behavior.

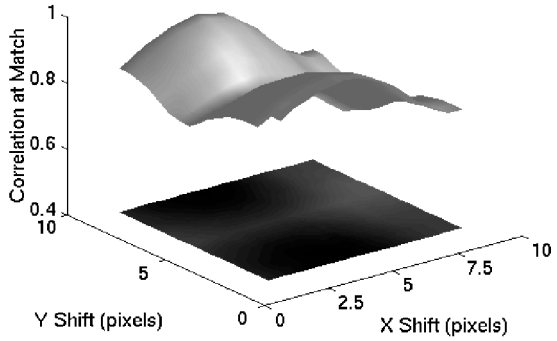

Fig. 8. Low-pass correlations for Haar wavelets at level 3 as a function of x and y phase shifts of the template through 64 possible template phases.

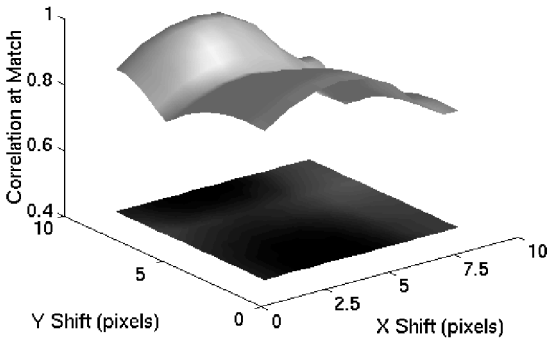

Fig. 9. Low-pass correlations for Daubechies wavelets at level 3 as a function of x and y phase shifts of the template through 64 possible template phases.

In Fig. 8, the minimum correlation peak with a value of 0.7032 occurs at an offset (4, 5) with respect to the peak. The mean for the Haar over the 64 offsets in Fig. 8 is 0.8823. The minimum correlation peak for the Daubechies is 0.7085 at an offset of (4, 5). The average peak height for the Daubechies wavelet is 0.8921, which is very close to the mean peak height for the Haar wavelet. In all phases, the globally maximum correlation value is at the correct registration position for both the Haar and Daubechies wavelets.

To compare this with the model of Section 2, consider the graphs of the 2D-model in Fig. 3. Average peak heights for the model exceed 0.96 for large features. The small-feature correlation peaks vary around 0.8. These two sets of correlations are, respectively, higher and lower than our observations for the real image and the results are consistent with an image being composed of a mix of large and small features. The worst-case peak heights of the model vary between about 0.2 and 0.9, depending on pulse width as compared to about 0.7 for the image data. This is also consistent. The correlation peaks compared here are maxima of signed correlations.

The 3D graphs in both Figs. 8 and 9 are remarkably flat, thereby indicating excellent robustness with respect to translation in real images. A somewhat different situation appears in Figs. 10 and 11, which show the correlations of high-pass wavelet coefficients. The wavelet correlations in these figures have an ideal maximum value of 2.000 because they sum the absolute values of high-pass correlations in two directions. Absolute values should be used instead of signed values when there exists the possibility of intensity inversion from image to image.

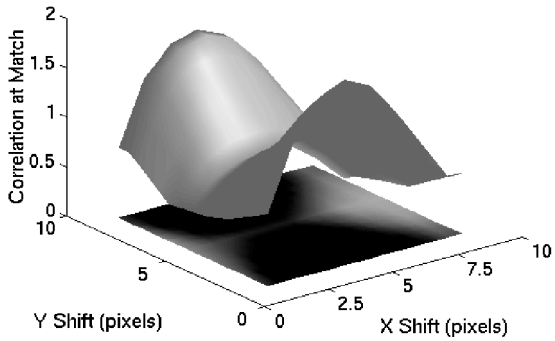

Fig. 10. Edge correlations for Haar wavelets at level 3 as a function of x and y phase shifts of the template through 64 possible template phases. The edge correlations are the sum of the absolute value of vertical and horizontal edge correlations.

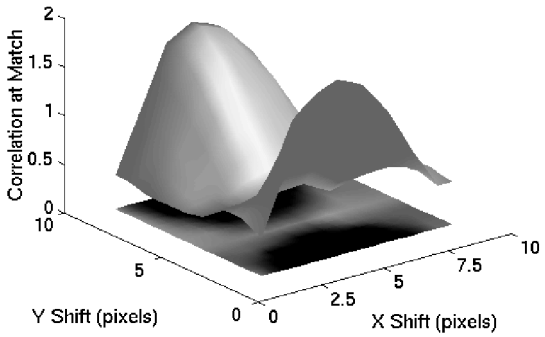

Fig. 11. Edge correlations for Daubechies wavelets at level 3 as a function of x and y phase shifts of the template through 64 possible template phases. The edge correlations are the sum of the absolute value of vertical and horizontal edge correlations.

For both Figs. 10 and 11, the correct registration point is usually at a local maximum of the correlation function and less often at a global maximum. The local search area was arbitrarily set to a square of size of 80×80 full-resolution pixels centered at the correct position. For the Haar filter, the correlation peak attains its upper bound 2.000 for a phase shift of 0. Of the 64 phases, for 42 and 59 of them the registration point is at a global and local maximum correlation, respectively.

The Daubechies high-pass correlations in Fig. 11 are somewhat poorer than the corresponding Haar correlations. The Daubechies coefficients depend on pixels that lie beyond the boundary of the template and extrapolation of existing pixel data provides the missing pixel data. The poorer correlations are due in part to the extrapolations. The data in Fig. 11 contain 31 and 45 correct global and local registrations, respectively, among the 64 phases.

To compare these results with the model of Section 2, we obtained minimum and average Haar correlation peaks of 0.22 and 0.50, respectively, in the low-high subband, and of 0.17 and 0.61 in the high-low subband. All values are very close to the data in Fig. 4. The comparable data for Daubechies are minimum and average correlation peaks of size 0.06 and 0.53, respectively, in the low-high subband and 0.09 and 0.49 in the high-low subband. This too is consistent with Fig. 4.

The comparison of Haar and Daubechies wavelets indicates about equal performance in the low-pass subband and slightly better performance for the Haar in the high-pass subband. Although both filters give imperfect results for the high-pass subband at this resolution level, the high-pass data can be used with caution.

Low correlation peaks characterized all cases for which the correct registration point was not at a local maximum. This suggests that one can register images with a high-pass subband correlation and discard the registration if the correlation fails to exceed a given threshold. In many of the cases for which the correct registration did not produce a local maximum, for all but four cases, it was within the top five peaks. The four are failures at Level 30.

Not shown are the high-pass results for Levels 1 and 2. At Level 1, $B = 4$, all global correlations are correct for both filters. For Level 2, the Haar global maxima are all correct and 13 of 16 Daubechies global maxima are correct. Level 3 is the lowest resolution possible for this particular template because, at lower resolution, the template contains too few pixels for a successful search. We conducted two other similar experiments using templates taken from different areas of the same image. The results of those experiments were consistent with the experiment reported above and are not reported separately.

4 Summary and Conclusions

The white-noise analysis, idealized mathematical model, and experimental data give consistent answers that wavelet registration of low-pass subbands is robust with respect to translation. The results satisfy the Sampling Theorem in that wavelets contain at least two samples of stable features. High-pass subbands exhibit translation sensitivity reported in the literature. The relative difference in performance of Haar and Daubechies filters is not significant.

We conclude that coarse-to-fine registration based on low-pass subbands should work well, provided that the subbands contain at least two samples of the smallest features of interest. High-pass subbands are usable with some caution, but are not reliable. In some phases, edge features can disappear in the high-pass subband and, thus, will not contribute to correlation.

This work leaves open a number of questions. Are there wavelet filters that substantially outperform Haar and Daubechies with respect to normalized correlation? Are there other approaches that capture edge information robustly at low-resolution? For Daubechies filters, how should data be extrapolated at boundaries? We hope to see these questions addressed in future work.

Acknowledgments

We wish to thank the anonymous referees for their helpful comments. This is a significantly enhanced version of the paper "Image Registration Using Wavelet Techniques" by H.S. Stone, J. Le Moigne, and M. McGuire which appeared in the *Proceedings of the 1997 AIPR Workshop*, October 1997.

References

[1] L. Brown, "A Survey of Image Registration Techniques," *ACM Computing Surveys*, vol. 24, no. 4, pp. 325-376, 1992.

[2] J. Le Moigne, W.J. Campbell, and R.F. Cromp, "An Automated Parallel Image Registration Technique of Multiple Source Remote Sensing Data," *IEEE Trans. Geoscience and Remote Sensing*, submitted.

[3] A.P. Cracknell and K. Paithoonwattanakij, "Pixel and Sub-Pixel Accuracy in Geometrical Correction of AVHRR Imagery," *Int'l J. Remote Sensing*, vol. 10, nos. 4-5, pp. 661-667, 1989.

[4] B.J. Devereux, R.M. Fuller, L. Carter, and R.J. Parsell, "Geometric Correction of Airborne Scanner Imagery by Matching Delaunay Triangles," *Int'l J. Remote Sensing*, vol. 11, no. 12, pp. 2,237-2,251, 1990.

[5] P.D. Fiore, "Image Registration Using Both Distance and Angle Information," *Proc. Int'l Conf. Image Processing*, vol. III, pp. 220-223, Oct. 1995.

[6] C.D. Kuglin and D.C. Hines, "The Phase Correlation Image Alignment Method," *Proc. IEEE 1975 Conf. Cybernetics and Society*, pp. 163-165, Sept. 1975.

[7] S. Alliney, G. Cortelazzo, and G.A. Mian, "On the Registrations of an Object Translating on a Static Background," *Pattern Recognition*, vol. 29, no. 1, pp. 131-141, Jan. 1996.

[8] B.S. Reddy and B.N. Chatterji, "An FFT-Based Technique for Translation, Rotation, and Scale-Invariant Image Registration," *IEEE Trans. Image Processing*, vol. 3, no. 8, pp. 1,266-1,270, Aug. 1996.

[9] P.E. Anuta, "Spatial Registration of Multispectral and Multitemporal Digital Imagery Using Fast-Fourier Transform Techniques," *IEEE Trans. Geoscience Electronics*, vol. 8, no. 4, pp. 353-368, Oct. 1970.

[10] J.W. Wong and E.L. Hall, "Scene Matching with Invariant Moments," *Computer Graphics and Image Processing*, vol. 8, pp. 16-24, 1978.

[11] D. Casasent, R. Schaefer, and R. Sturgill, "Optical Correlation Filter Fusion for Object Detection," *Applied Optics*, vol. 31, no. 29, pp. 6,255-6,263, Oct. 1992.

[12] D. Casasent, J.S. Smokelin, and R. Schaefer, "Optical Correlation Filter Fusion for Object Detection," *Optical Eng.*, vol. 33, no. 6, pp. 1,757-1,766, June 1994.

[13] J. Le Moigne, "Parallel Registration of Multi-Sensor Remotely Sensed Imagery Using Wavelet Coefficients," *Proc. SPIE O/E Aerospace Sensing, Wavelet Applications*, pp. 432-443, Apr. 1994.

[14] M. Khosravi and R. Schaefer, "Template Matching Based on a Grayscale Hit-or-Miss Transform," *IEEE Trans. Image Processing*, vol. 5, no. 6, pp. 1,060-1,066, June 1996.

[15] H.S. Stone, "Progressive Wavelet Correlation Using Fourier Methods," *IEEE Trans. Signal Processing*, vol. 47, no. 1, pp. 97-107, Jan. 1999.

[16] R.L. Allen, F.A. Kamangar, and E.M. Stokely, "Laplacian and Orthogonal Wavelet Pyramid Decompositions in Coarse-to-Fine Registration," *IEEE Trans. Signal Processing*, vol. 41, no. 12, pp. 3,536-3,541, Dec. 1993.

[17] J.P. Djamdji, A. Bijaoui, and R. Maniere, "Geometrical Registration of Images: The Multiresolution Approach," *Photogrammetric Eng. and Remote Sensing J.*, vol. 59, no. 5, pp. 645-653, 1993.
[18] Q. Zheng and R. Chellappa, "A Computational Vision Approach to Image Registration," *IEEE Trans. Image Processing*, vol. 2, no. 3, pp. 311-326, 1993.
[19] M. Corvi and G. Nicchiotti, "Multiresolution Image Registration," *Proc. 1995 IEEE Int'l Conf. Image Processing*, pp. 224-227, 1995.
[20] J. Le Moigne, "Towards a Parallel Registration of Multiple Resolution Remote Sensing Data," *Proc. 1995 Int'l Geoscience and Remote Sensing Symp.*, pp. 1,011-1,013, July 1995.
[21] H.H. Li and Y.-T. Zhou, "A Wavelet-Based Point Feature Extractor for Multi-Sensor Image Registration," *Proc. SPIE Aerosense Wavelet Applications III*, pp. 524-534, Apr. 1996.
[22] E.P. Simoncelli, W.T. Freeman, E.H. Adelson, and D.J. Heeger, "Shiftable Multiscale Transforms," *IEEE Trans. Information Theory*, vol. 38, no. 2, pp. 587-607, Mar. 1992.
[23] D. Casasent and R. Shenoy, "New Gabor Wavelets with Shift-Invariance for Improved Time-Frequency Analysis and Signal Detection," *Proc. Wavelet Applications III*, vol. 2,762, pp. 244-255, Apr. 1996.
[24] N. Saito and G. Beylkin, "Multiresolution Representations Using the Autocorrelation Functions of Compactly Supported Wavelets," *IEEE Trans. Signal Processing*, vol. 41, no. 12, pp. 3,584-3,590, Dec. 1993.
[25] I. Cohen, S. Raz, and D. Malah, "Shift Invariant Wavelet Packet Bases," *Proc. 1995 Int'l Conf. Acoustics, Speech, and Signal Processing*, pp. 1,081-1,084, May 1995.
[26] J. Liang and T.W. Parks, "Translation Invariant Wavelet Transforms with Symmetric Extensions," *Proc. 1996 IEEE Digital Signal Processing Workshop*, pp. 69-72, Sept. 1996.
[27] G. Strang and T. Nguyen, *Wavelets and Filter Banks*. Wellesley, Mass.: Wellesley-Cambridge Press, 1996.