Advanced discrete wavelet transform fusion algorithm and its optimization by using the metric of image quality index

Yufeng Zheng
University of Louisville
Department of Psychological and Brain Sciences
317 Life Sciences Bldg
Louisville, Kentucky 40292
E-mail: yufeng.zheng@louisville.edu

Edward A. Essock
University of Louisville
Department of Psychological & Brain Sciences
Louisville, Kentucky 40292
and
University of Louisville
Department of Ophthalmology & Visual Science
Louisville, Kentucky 40292

Bruce C. Hansen
University of Louisville
Department of Psychological & Brain Sciences
Louisville, Kentucky 40292

Abstract. Of the many image fusion methods, the discrete wavelet transform (DWT) and various pyramids (e.g., the Laplacian pyramid) are among the most common and effective. For quantitative evaluation of the quality of fused imagery, the root mean square error (RMSE) is the most reasonable measure of quality if a ''ground truth'' image is available; otherwise, the entropy, spatial frequency, or image quality index (IQI) can be calculated and evaluated. Here, an advanced discrete wavelet transform (*a*DWT) method that incorporates principal component analysis (PCA) and morphological processing into a regular DWT fusion algorithm is presented. Specifically, a principle vector is derived from two input images and then applied to two of the images' approximation coefficients at the highest DWT transform scale. For the detail coefficients at each transform scale, the larger absolute values are chosen and subjected to a neighborhood morphological processing procedure that serves to verify the selected pixels by using a ''filling'' and ''cleaning'' operation. Furthermore, the *a*DWT has two adjustable parameters—the number of DWT decomposition levels and the length of the selected wavelet that determinately affect the fusion result. An iterative fusion process that was optimized with the established metric—IQI—is then implemented. Experimental results tested on four types of inhomogeneous imagery show that the iterative *a*DWT achieves the best fusion compared to the pyramid or the DWT methods judged on both the IQI metric and visual inspection. [DOI: 10.1117/1.1871812]

Subject terms: image fusion; image pyramid; image quality index; iterative fusion; spatial frequency; wavelet transform.

Paper 040306 received May 21, 2004; revised manuscript received Aug. 24, 2004; accepted for publication Sep. 22, 2004; published online Mar. 23, 2004.

1 Introduction

Along with the development of new imaging sensors arises the requirement for a meaningful combination of all employed imaging sources (i.e., image fusion). There are many applications of image fusion, some of which include medical imaging, remote sensing, nighttime operations, and multispectral imaging. A general definition of image fusion has been stated as ''the combination of two or more different images to form a new image by using a certain algorithm.''[1] The actual fusion process can take place at different levels of information representation. A generic categorization of these different levels can be sorted in ascending order of abstraction: pixel, feature, and decision levels. This paper focuses on the so-called pixel-level fusion process, whereby a composite image must be built of multiple (typically two) input images.

Image fusion is a tool that serves to combine multisource imagery by using advanced image processing techniques. Specifically, it aims at the integration of disparate and complementary data to enhance the information apparent in the images, as well as to increase the reliability of the interpretation—ultimately leading to more accurate data[2] and increased utility.[3–4] In addition, it has been shown that fused data can provide for robust operational performance of human users achieved through increased confidence, reduced ambiguity, improved reliability, and improved classification.[3,5]

A general framework of image fusion can be found in Ref. 6. In pixel-level image fusion, some general requirements[7] are imposed on the fusion result: (1) the fusion process should preserve all relevant information of the input imagery in the composite image (pattern conservation) and (2) the fusion scheme should not introduce any artifacts or inconsistencies that would distract the human observer or disrupt subsequent processing stages. Consequently, quantitative evaluation of the quality of fused imagery is considered very important for an objective comparison of different types of fusion algorithm performance. In addition, a quantitative metric can potentially be used as feedback to the fusion algorithm to further improve fused image quality.

Driven by the wide range of applications for image fusion, many fusion algorithms have been developed—two common methods being the discrete wavelet transform[8–11] (DWT) and various pyramids[12,13] (such as the Laplacian, contrast, gradient, and morphological pyramids)—whereas

only a few metrics are available for quantitative evaluation of the quality of fused imagery. For example, the root mean square error (RMSE) is the most natural measure of image quality if a ''ground truth'' image is available; however, for realistic image fusion problems, there are no ground truths. The spatial frequency (SF) or entropy have been used to measure how much information is contained in an image and have been subsequently used in the fusion rules (i.e., the rules that decide which input should be selected for the fused image),[14] but they cannot directly measure the fused image quality. Beauchemin et al. presented a method that uses local variance for image fusion assessment[15] that still requires a comparison with the measurement of ground truth. Leung et al. proposed the image noise index (INI) based on entropy calculation to measure the fused image quality.[16] However, this method requires the exact reverse process of an image fusion procedure, which is almost impractical to be realized for most fusion processes such as DWT or pyramid fusion methods. Piella and Heijmans recently presented a new metric for image fusion, the image quality index[17] (IQI), which measures how similar the fused image is to both input images by evaluating their correlation and distortion. The values of IQI are within [0,1] regardless of image content, where the ideal value is 1 (if two input images are identical). In addition, the IQI value is easily computed and fairly accurate (compared to other available metrics). Therefore, this IQI metric can be used to direct an iterative fusion process that updates its parameters based on the IQI value, thus enabling the best parameter settings (thereby producing the ideal fused version of the input images) for different imagery.

Besides the various pyramid methods, an advanced wavelet transform (*a*DWT) method that incorporates principal component analysis (PCA) and morphological processing into a regular DWT fusion algorithm was recently presented.[18] Furthermore, experimental results showed an important relationship between the fused image quality and the wavelet properties:[18] a higher level of DWT decomposition (with smaller image resolution at higher scale) or a lower order of wavelets (with shorter length) usually resulted in a more sharpened fused image. This means that we can use the level of DWT decomposition and the length of a wavelet as the parameters of an iterative DWT-based fusion algorithm. Together with the IQI metric, an iterative *a*DWT can be realized.

The subsequent sections of this paper are organized as follows. The currently used metrics—RMSE, IQI, and SF—are presented in Sec. 2, where the extension of SF is developed. Next is a full description of pyramid methods and the *a*DWT followed by the introduction of the iterative *a*DWT that is directed by the IQI metric. Last, the experimental results and discussion are presented, followed by conclusions.

2 Image Quality Metrics

As mentioned in the introduction, the general requirements of an image fusion process are that it should preserve all valid and useful pattern information from the source images, while at the same time not introducing artifacts that could interfere with subsequent analyses.[19] However, it is almost impossible to combine images without introducing some form of distortion. In the current body of fusion literature, image fusion results tend to be evaluated by inspection (i.e., visually) or objectively (e.g., by either human performance or image-quality measures). Here, we focus on quantitative image-quality metrics that can be carried out automatically by computers. Three commonly used measures are described in the following, one of which, RMSE, requires a reference (ground truth) image, while the others do not. Note, however, that the best criterion should be linked with the specific application.

2.1 *RMSE*

The RSME between the reference image (ground truth) and the fused image is

$$\mathrm{RMSE}=\left\{\frac{\Sigma_{i=1}^{M}\Sigma_{j=1}^{N}[I_R(i,j)-I_F(i,j)]^2}{M\times N}\right\}^{1/2}, \tag{1}$$

where $I_R(i,j)$ and $I_F(i,j)$ are the image pixel values of the reference image and the fused image, respectively, and $M\times N$ is the image size.

2.2 *Image Quality Index*

The IQI was recently introduced by Wang and Bovik.[20] Given two sequences $\mathbf{x}=(x_1,\ldots,x_n)$ and $\mathbf{y}=(y_1,\ldots,y_n)$, let $\bar{x}$ denote the mean of $\mathbf{x}$, and let σ_x and σ_{xy} denote the variance of $\mathbf{x}$ and covariance of $\mathbf{x}$ and $\mathbf{y}$, respectively. The global quality index of two vectors is defined as

$$Q_0(x,y)=\frac{4\sigma_{xy}\overline{\mathbf{xy}}}{(\overline{x}^2+\overline{y}^2)(\sigma_x^2+\sigma_y^2)}, \tag{2}$$

which can be decomposed as

$$Q_0(x,y)=\frac{\sigma_{xy}}{\sigma_x\sigma_y}\,\frac{2\overline{\mathbf{xy}}}{\overline{x}^2+\overline{y}^2}\,\frac{2\sigma_x\sigma_y}{\sigma_x^2+\sigma_y^2}. \tag{3}$$

Note that the first component in Eq. (3) is the correlation coefficient between $\mathbf{x}$ and $\mathbf{y}$. This value is a measure for the similarity of the vectors $\mathbf{x}$ and $\mathbf{y}$, and takes on values between -1 and 1. Keep in mind that in this case (image quality evaluation), the values x_i and y_i are positive gray-scale values. The second component in Eq. (3) corresponds to the luminance distortion which has a dynamic range of [0,1]. The third factor in Eq. (3) measures the contrast distortion, its range is also [0,1]. In summary, $Q_0\in[0,1]$; and the maximum value $Q_0=1$ is achieved when $\mathbf{x}$ and $\mathbf{y}$ are identical.

Piella and Heijmans introduced a weighting procedure into the Q_0 calculation,[17] where $S(I_A)$ denotes the ''saliency'' of image *A*. It should reflect the local relevance of image *A* that can depend on local variance, contrast, sharpness, or entropy. Given the local saliencies of two input images, *A* and *B*, we compute a local weight λ indicating the relative importance of image *A* compared to image *B*: the larger λ, the more weight is given to image *A*. A typical choice for λ is

$$\lambda=S(I_A)/[S(I_A)+S(I_B)]. \tag{4}$$

Then, the fused image quality measure (i.e., the weighted image quality index) can be defined as

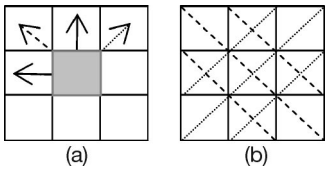

Fig. 1 Illustration of SF calculation where each small square denotes a pixel: (a) four-directional SFs with arrows pointed to the gradient directions and (b) two diagonal directions; where the dashed lines stand for main diagonal directions and the dotted lines represent secondary diagonal directions.

$$Q_w = \lambda Q_0(I_A, I_F) + (1-\lambda) Q_0(I_B, I_F). \tag{5}$$

Since image signals are generally nonstationary, it is more appropriate to measure the weighted image quality index Q_w over local regions (e.g., by parsing the entire image into a set of ''blocks'') and then combine the different results into a single measure. Piella and Heijmans also suggested using local variance as the saliency of an image, i.e., $S(I_A) = \sigma_A$.

2.3 *Spatial Frequency*

SF is used to measure the overall activity level of an image,[21] which is defined as follows.

$$\mathrm{SF} = [(\mathrm{RF})^2 + (\mathrm{CF})^2]^{1/2}, \tag{6}$$

where RF and CF are row frequency and column frequency, respectively.

$$\mathrm{RF} = \left\{ \frac{1}{MN} \sum_{i=1}^{M} \sum_{j=2}^{N} [I(i,j) - I(i,j-1)]^2 \right\}^{1/2}, \tag{7a}$$

$$\mathrm{CF} = \left\{ \frac{1}{MN} \sum_{j=1}^{N} \sum_{i=2}^{M} [I(i,j) - I(i-1,j)]^2 \right\}^{1/2}. \tag{7b}$$

Similar to the definitions of RF and CF, spatial frequency along two diagonal directions [see Fig. 1(b)], termed as main diagonal SF (MDF) and secondary diagonal SF (SDF), can be defined as

$$\mathrm{MDF} = \left\{ w_d \frac{1}{MN} \sum_{i=2}^{M} \sum_{j=2}^{N} [I(i,j) - I(i-1,j-1)]^2 \right\}^{1/2}, \tag{8a}$$

$$\mathrm{SDF} = \left\{ w_d \frac{1}{MN} \sum_{j=1}^{N-1} \sum_{i=2}^{M} [I(i,j) - I(i-1,j+1)]^2 \right\}^{1/2}, \tag{8b}$$

where $w_d = 1/\sqrt{2}$ is a distance weight; similarly it can be considered that $w_d = 1$ in Eqs. (7a) and (7b).

Then the overall SF of an image becomes

$$\mathrm{SF} = [(\mathrm{RF})^2 + (\mathrm{CF})^2 + (\mathrm{MDF})^2 + (\mathrm{SDF})^2]^{1/2}, \tag{9}$$

which is a combination of four directional SFs [see Fig. 1(a)]. With Eq. (9) we can calculate and compare the SFs of input images or the fused image (SF_F).

2.4 *Entropy*

$$\mathrm{entropy} = -\sum_{l=0}^{L-1} p(l) \log_2 p(l), \tag{10}$$

where $p(l)$ is the probability of gray level l (which can be computed by analyzing the image histogram), and the dynamic range of the analyzed image is $[0, L-1]$ (usually, $\mathrm{L} = 256$).

3 Image Fusion Methods

3.1 *Pyramid Methods*

Image pyramids were initially described for a multiresolution image analysis and as a model for the binocular fusion in human vision.[19,22] An image pyramid can be described as collection of low- or bandpass copies of an original image in which both the band limit and sample density are reduced in regular steps. The basic strategy of image fusion based on pyramids is to use a feature selection rule to construct a fused pyramid representation from the pyramid representations of the original images. The composite image is obtained by taking the inverse pyramid transform. Several pyramid-based fusion schemes have been proposed recently,[22–33] and they are briefly reviewed next.

3.1.1 *Laplacian pyramid*

The Laplacian pyramid was first introduced as a model for binocular fusion in human stereo vision,[19,22] where the implementation used a Laplacian pyramid and a maximum selection rule at each point of the pyramid transform. Essentially, the procedure involves a set of bandpass copies of an image and is referred to as the Laplacian pyramid due to its similarity to a Laplacian operator. Each level of the Laplacian pyramid is recursively constructed from its lower level by applying the following four basic steps: blurring (low-pass filtering); subsampling (reduce size); interpolation (expand); and differencing[23] (to subtract two images pixel by pixel). In the Laplacian pyramid, the lowest level of the pyramid is constructed from the original image.

3.1.2 *Ratio of low-pass pyramid and contrast pyramid*

The ratio of low-pass (RoLP) or contrast pyramid, which was introduced by Toet[24] and Toet et al.,[25] is very similar to a Laplacian pyramid. Originally the RoLP pyramid was explicitly intended for use by human observers. Every level of the RoLP pyramid is the ratio of two successive levels of the Gaussian pyramid. In Refs. 24 to 27, an RoLP pyramid and the maximum selection rule were used for visible-to-IR image fusion.

3.1.3 *Gradient pyramid*

The gradient pyramid can be generated by applying gradient operators to each level of the Gaussian pyramid.[28] This produces horizontal, vertical, and diagonal pyramid sets for each source in the Gaussian pyramid. Burt and Lolczynski[29] proposed an image fusion scheme that was based on a gradient pyramid and an activity measure within a small window rather than just a single point. In Ref. 30,

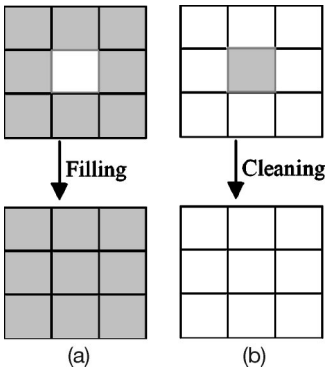

Fig. 2 Morphological processing (of a 3×3 region): (a) filling and (b) cleaning. Shaded pixels were chosen from image *A*, and white pixels were selected from image *B*.

Richard et al. demonstrated the qualitative and quantitative results of the preceding three image data fusion algorithms and their target signature variation.

3.1.4 *Morphological pyramid*

A morphological pyramid can be constructed by the successive filtering of the original image with a sequence of nonlinear morphological operators (such as the open-close filter) and a specialized subsampling scheme.[31] The application of morphological pyramid to image fusion can be referenced in Refs. 32 and 33.

3.2 *DWT and a*DWT *Methods*

As with a pyramid method, the regular DWT method is also a multiscale analysis method. In a regular DWT fusion process, DWT coefficients from two input images are fused (pixel by pixel) by choosing the average of the approximation coefficients at the highest transform scale and the larger absolute value of the detail coefficients at each transform scale. Then an inverse DWT is performed to obtain the fused image. In the *a*DWT method, PCA (Ref. 34) and morphological processing are incorporated into the DWT fusion algorithm. Specifically, at each DWT scale of a particular image, the DWT coefficients of a 2-D image consist of four parts: approximation, horizontal detail, vertical detail, and diagonal detail. We then apply PCA to the two input images' approximation coefficients at the highest transform scale, that is, we fuse them by using the principal eigenvector (corresponding to the larger eigenvalue) derived from the two original images, as described by

$$C_F=(a_1C_A+a_2C_B)/(a_1+a_2), \tag{11}$$

where C_A and C_B are approximation coefficients trans-

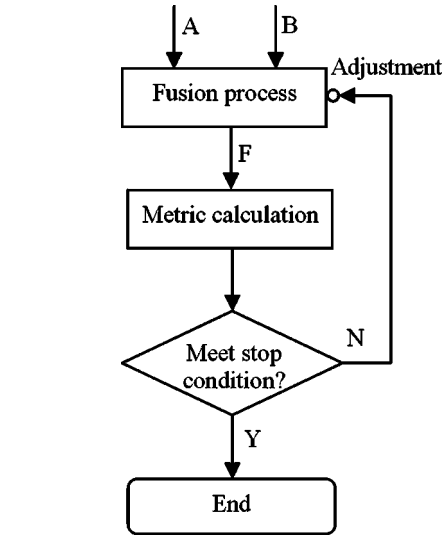

Fig. 3 Diagram of a general iterative back-propagation fusion process: *A* and *B* are input images and *F* is the fused image.

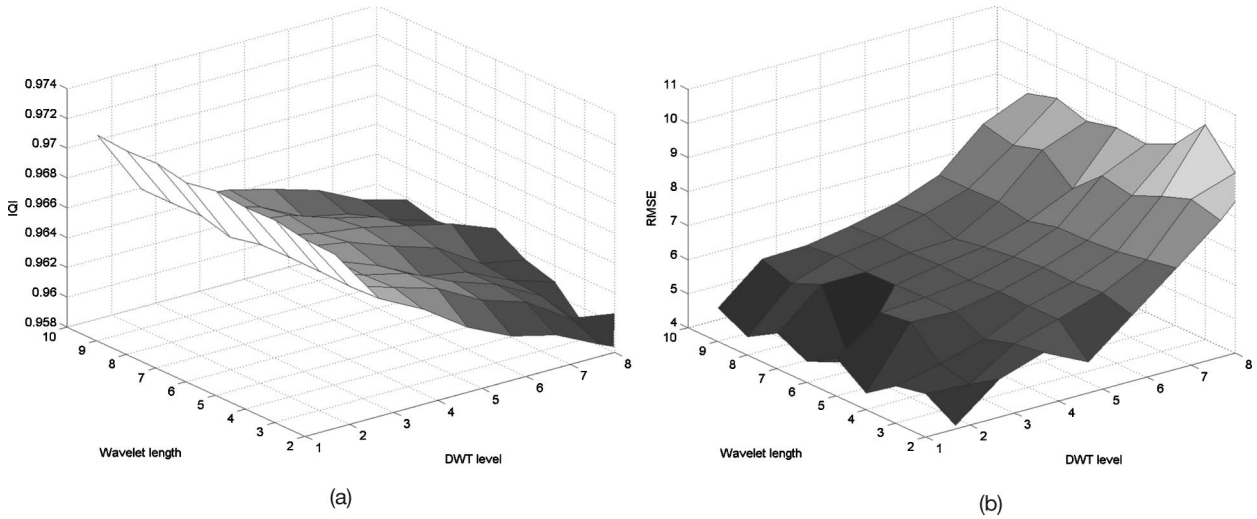

Fig. 4 Fusion metric distributions with varying parameters of DWT level (L_d) and wavelet length (L_w) while fusing the simulated pair of "Lena-I2" for (a) distribution of IQI (L_d, L_w) and (b) distribution of RMSE (L_d, L_w). (a) Shows a monotonically decreasing function of IQI (L_d). By comparing (a) and (b), we see that two metrics (IQI and RMSE) are quite consistent, which means that a large IQI corresponds a small RMSE.

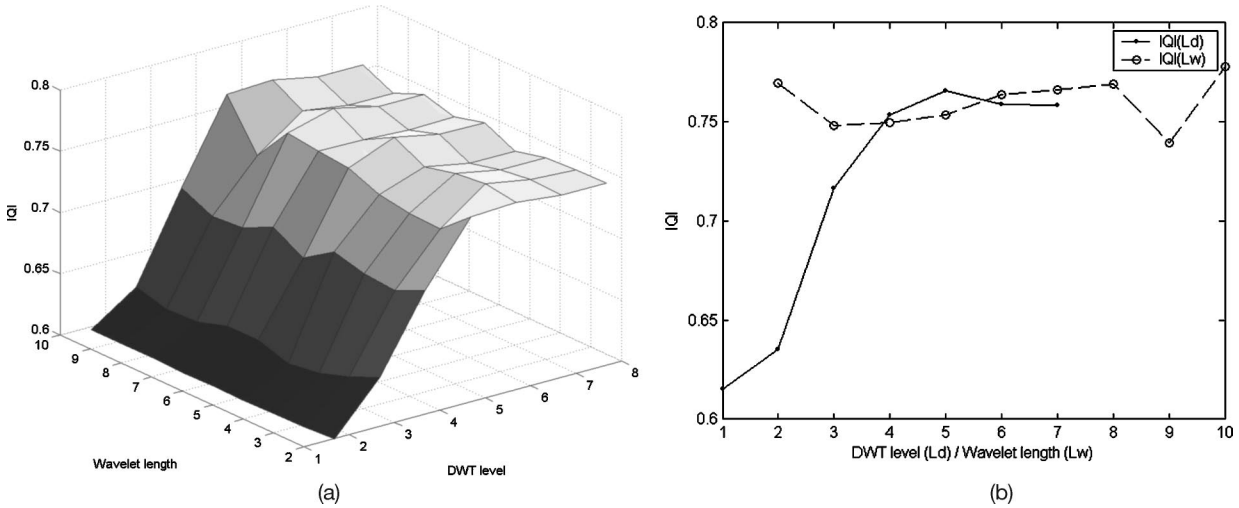

Fig. 5 IQI distribution with varying parameters of DWT level (L_d) and wavelet length (L_w) while fusing the medical image pair for (a) distribution of IQI (L_d, L_w) and (b) distribution of IQI (L_d) and IQI (L_w)—two slices of (a) at $L_w=4$ and $L_d=4$, respectively. (b) Clearly shows an almost increasing function of IQI (L_d), while IQI (L_w) is locally fluctuating.

formed from input images A and B; C_F represents the fused coefficients; and a_1 and a_2 are the elements of the principal eigenvector, which are computed by analyzing the original input images. (Note we do not analyze C_A and C_B alone because their sizes at the highest transform scale are too small to conduct an accurate result.) In practice, we have to convert 2-D images into 1-D vectors by simply stacking each image column by column so that the principle component can be computed. Note that the denominator in Eq. (11) is used for normalization so that the fused image has the same energy distribution as the original input images.

For the detail coefficients (the other three quarters of the coefficients) at each transform scale, the larger absolute values are selected, followed by neighborhood morphological processing, which serves to verify the selected pixels by using a ''filling'' and ''cleaning'' operation (i.e., the operation fills or removes isolated pixels locally, as shown in Fig. 2). For example, in the fused-coefficients image, if the detail coefficient of the central pixel is selected from input image B, but the other eight surrounding coefficients are selected from input image A [as shown in Fig. 2(a)], then a filling operation will replace this central coefficient (from image B) with the one from image A. The filling operation takes no action when the number of the surrounding eight pixels had less than eight coefficients selected from image A. Such an operation (similar to smoothing) can increase the consistency of coefficient selection, thereby reducing the distortion in the fused image.

Note that the highest DWT level ($L_{d\,\max}$) that can be decomposed depends on the input image size. However, the size of smallest transformed image should not be less than 2×2. Thus, we have

$$L_{d\,\max}=\mathrm{int}\{\log_2[\min(M,N)/\min(m_0,n_0)]\}, \tag{12}$$

where the image size is $M\times N$, $m_0\times n_0$ represents the size of transformed image by DWT at the highest scale, min means taking the minimal value, and int signifies taking the integral part.

3.3 *Iterative a*DWT *Directed by the IQI*

A general iterative model for any image fusion procedure is proposed and is shown in Fig. 3. Here the IQI value is calculated to measure the fused image quality by the aDWT, then it is fed back to the fusion algorithm (aDWT) to achieve a better fusion by directing parameter adjustment. To make this model work properly, some requirements are imposed: (1) the fusion process must have some adjusted parameters that definitely affect the metric value;

Fig. 6 Image fusion with the simulated pair from image ''Lena-I2.'' The original image size is 512×512 pixels. All images were scaled to fit in the paper space. (a) Lena image (reference or ground truth), (b) input image A, (c) input image B, (d) image fused by iterative aDWTi, (e) image fused by Laplacian pyramid, and (f) image fused by regular DWT.

Table 1 The fused image quality evaluations of four metrics (entropy, SF, IQI, and RMSE).

Imagery	Ref.	Input A	Input B	Metric	Lapl.	Ratio	Contr.	Grad.	Morph.	DWT	*a*DWT	*a*DWTi
"Lena-I1" (1,4) (N_i=5)	7.312	7.289	7.287	Entropy	7.457	7.378	7.441	7.416	7.353	7.427	7.443	7.442
	20.979	16.375	18.621	SF	20.839	16.629	2.909	19.176	19.152	20.680	21.193	18.647
				IQI	*0.976*	0.974	0.975	0.960	0.965	0.975	0.975	**0.979**
				RMSE	3.048	5.997	**1.640**	13.908	8.760	*1.853*	2.713	2.646
"Lena-I2" (1,4) (N_i=5)	7.312	7.365	7.30	Entropy	7.532	7.358	7.422	7.567	7.408	7.530	7.470	7.450
	20.979	16.703	18.936	SF	21.107	15.293	18.518	19.890	19.673	21.306	21.713	18.886
				IQI	0.962	0.958	0.962	0.952	0.955	0.960	*0.966*	**0.972**
				RMSE	12.783	15.566	17.542	13.054	9.652	12.489	*6.378*	**5.086**

Eight fusion methods (five pyramids plus three DWTs) were tested on two simulated image pairs—"Lena-I1" and "Lena-I2." The converged parameters of *a*DWTi (iterative *a*WT)—the DWT level and the wavelet length (L_d,L_w) and the number undergone iterations (N_i) are listed below the imagery name in the first column. All seven noniterative fusion methods (expect for *a*DWTi) were done with default parameter (L_d,L_w)=(4,4). The entropies and spatial frequencies of the reference (ground truth) image and input images are also shown in the table.

(2) there must be an ideal value for the metric that is used as a "destination" for judging the iteration convergence; and (3) the error of the metric values (obtained by subtracting the current metric value from the ideal metric value) should indicate the error direction that is used to direct the parameter adjustment (i.e., increment or decrement). Requirements 1 and 2 are necessary, while requirement 3 is optional but will expedite the convergence of iteration.

Previous experiments[18] have pointed to an important relationship between the fused image quality and the wavelet properties, that is, a higher level DWT decomposition (with smaller image resolution at higher scale) or a lower order of wavelets (with a shorter length) usually resulted in a more sharpened fused image. If measured with IQI, the value usually tends to be large for a sharpened image. This means that we can use the level of DWT decomposition (L_d) and the length of a wavelet (L_w) as adjusted parameters of an iterative *a*DWT algorithm. With the definition of IQI, we know that it has an ideal value, 1, but it can not give the error direction because $0<\mathrm{IQI}\leqslant 1$.

Of course, a termination condition is required to stop the fusion iteration (Fig. 3). The fusion iteration is designed to stop when one of the following conditions is met: (1) convergence at the destination point—the absolute value of (IQI-1), i.e., abs(IQI-1), is smaller than a designated tolerance error (for example, 0.005); (2) convergence at some point but not at the destination point—that is, when there is no significant change of the IQI value between two adjacent iterations (e.g., the change <0.0005); (3) no convergence resulting from an IQI value that is decreasing or fluctuating but generally decreasing for subsequent iterations; and (4) no convergence resulting when the parameter boundaries are reached. In implementing the iteration of a fusion procedure, appropriate boundaries of varying parameters should be designated based on the definition of parameters and the context. The details of implementation are depicted in the next section.

4 Experimental Results and Discussion

The experiments were organized to test a total of eight image fusion algorithms (i.e., five pyramid-based and three DWT-based methods) on four types of imagery (already well registered and with 256 gray levels): (1) two simulated image pairs, (2) three frequently used samples, (3) image intensified and IR paired images [obtained from the Army Night Vision Electro-optic Sensors Directorate (NVESD)], and (4) near-IR and IR paired images (obtained by our lab). We applied each fusion method to the four imagery groups, and calculated their fusion performance with the four quantitative measures. All experimental results are given in

Table 2 The fused image quality evaluations of three metrics (entropy, SF and IQI).

Imagery	Input A	Input B	Metric	Lapl.	Ratio	Contr.	Grad.	Morph.	DWT	*a*DWT	*a*DWTi
Clock (6,7) (N_i=15)	6.978	6.924	Entropy	7.140	7.309	7.227	7.356	6.965	7.221	7.461	7.409
	14.071	11.171	SF	13.235	11.033	16.095	12.222	12.150	14.238	16.339	16.540
			IQI	0.889	0.808	**0.928**	0.876	0.855	0.894	0.924	*0.927*
Medical (5,10) (N_i=14)	1.713	5.656	Entropy	6.468	3.595	1.741	6.455	5.738	5.689	6.787	6.874
	21.334	20.858	SF	19.452	7.234	16.134	19.625	20.712	12.913	27.811	27.918
			IQI	0.698	0.163	0.300	0689	0.600	0.425	*0.750*	**0.780**
Remote (5,3) (N_i=14)	5.445	5.454	Entropy	7.272	6.996	6.902	7.597	6.616	7.091	7.370	7.251
	53.769	46.857	SF	42.468	28.796	40.078	42.297	31.709	35.409	60.839	60.916
			IQI	0.831	0.635	0.77	0.788	0.697	0.739	*0.866*	**0.868**

Eight fusion methods were tested on three commonly used samples. The converged parameters of *a*DWTi−(L_d,L_w) and N_i are listed below the imagery name in the first column. All seven noniterative fusion methods were done with default parameters (L_d,L_w)=(4,4). The entropies and spatial frequencies of input images are also shown in the table.

Tables 1 to 4 in Sec. 4.1 where the best (bold) and second best (italic) performances are highlighted. Some selected fused images are illustrated in Figs. 6 to 13 in Sec. 4.1.

For imagery where the ground truth image is available, the RMSE is the most reasonable measure that indicates how close this fused image is to the ground truth. For imagery without ground truths, however, the IQI is relatively reliable, which gives a value (between 0 and 1) as an indicator of how similar the fused image is to both input images. Since the potential artificial distortion during the fusion process can also increase the entropy or SF values of the fused image, no judgment can be made by directly comparing the entropy or SF values (refer to Table 1 in Sec. 4.1). However, they still show how much information was contained in that image.

The DWT decomposition level, $L_d=4$ (used more often in the image fusion literature), was designated for all seven fusion algorithms except for the iterative *a*DWT (*a*DWTi). In the regular DWT method, the second-order Daubechies wavelets (a conventional choice) were used, while the fourth-order symlets (as modifications to the Daubechies family that are nearly symmetrical wavelets) were selected for the *a*DWT. In the *a*DWTi, symlets were also selected with adjustable parameters of the DWT decomposition level (L_d) and the wavelet length (L_w). The extreme values of two parameters are given as $1 \leqslant L_d \leqslant L_{d\,\max}$ and $2 \leqslant L_w \leqslant 10$. Experiments showed that the function IQI (L_d) is approximately monotonic (i.e., either increasing or decreasing) while IQI (L_w) changes in a less predictable fashion (see Figs. 4 and 5). In addition, varying L_d has a larger impact on the fused image quality than changing L_w. Therefore, in the iterative *a*DWT algorithm, parameters L_d and L_w were adjusted separately, which dramatically reduced the computation loads (refer to the number of iterations N_i in Table 1 in Sec. 4.1) comparing with jointly varying parameters (for instance, $N_i=8\times 9=72$ for a 512×512 image) achieved with the following steps: (1) starting with the given initializations of $L_d=1$ and $L_w=4$; (2) next was to vary L_d until IQI decreases; (3) with this fixed L_d, change L_w from 2 to 10; (4) last, find the maximum IQI. The current experiments showed that this separate parameter changing works perfectly to find the optimal IQI (Table 2 in Sec. 4.1).

4.1 *Experimental Results*

In the following analysis, all comparisons focused on the IQI metric because there is no standard or reference value for the entropy or SF metric, however, they are also given in all tables. To save space, five pyramid fusions were abbreviated as: Lapl. (Laplacian), Ratio, Contr. (contrast), Grad. (gradient), and Morph. (morphological).

4.1.1 *Two simulated image pairs (from "Lena")*

Two input images were generated by filtering a given image ["Lena" image, the ground truth or reference image, as shown in Fig. 6(a)] with a 5×5 Gaussian window. The "Lena" image was divided into 16 blocks, eight of which were randomly chosen to be blurred as one of the input images, referred to as image A. The other eight blocks were blurred to form image B. This resulted in "Lena" imagery 1 ("Lena-I1"). A second set, "Lena" imagery 2 ("Lena-I2"), was constructed by further reducing 16 gray levels after Gaussian blurring. The orthogonal stripes, as shown in Figs. 6(b) and 6(c), were caused by sliding the window along image block borders. It is more difficult to fuse "Lena-I2" perfectly because the input images differ in terms of brightness and contrast. The image quality evaluation results of the fused images by the eight algorithms are given in Table 1 where the converged parameters (L_d,L_w) and the number of iterations (N_i) for the *a*DWTi are also shown (below the imagery name). For "Lena-I1," judged on RMSE, the best method is the contrast pyramid, however, *a*DWTi did the best according to the measures of IQI. For "Lena-I2," the *a*DWTi had the best performance with both measures of RMSE and IQI, and the *a*DWT did the second best based on the IQI measure. Only three fused images (same as following figure illustrations) were illustrated in Figs. 6(d) to 6(f) because the Laplacian pyramid was shown to be the best among five pyramid methods, and the regular DWT was compared to the *a*DWTi. The fused image by the *a*DWT was not shown here since it is quite close to that of *a*DWTi (refer to the IQI measures in Table 1). In fact, it is quite hard to visually examine the difference between these fused images, especially from the scaled hardcopies. (That is particularly true for "Lena-I1," which is not displayed here.)

Figure 4 shows two function distributions of IQI (L_d,L_w) and RMSE (L_d,L_w) while fusing the simulated pair of "Lena-I2." Figure 4(a) clearly shows that IQI (L_d) is a monotonically decreasing function. By comparing Figs. 4(a) and 4(b), we see that two metrics (IQI and RMSE) are very consistent, which means that a large IQI corresponds to a small RMSE. The consistency between IQI and RMSE was also shown in Table 1, where for "Lena-I2," the algorithm of smallest RMSE coincided with that of largest IQI. In "Lena-I1," the IQI corresponding to the smallest RMSE (1.640) is 0.975, which is very close the largest IQI, 0.979 (where RMSE=2.646).

4.1.2 *Three frequently used samples*

Three commonly used image samples that were tested included two clocks imaged at different focal planes; two different types of medical imagery [computed tomography (CT) and magnetic resonance imaging (MRI)] and a pair of remote sensing images (IR and low-light sensors). The quantitative assessments of the fused images are listed in Table 2. For the medical and remote pairs, the *a*DWTi's performance was the best according to the metric of IQI, and the *a*DWT was the second best. For the clock pair, the contrast pyramid did the best with IQI (Contr.)=0.928, however, the IQI (*a*DWTi)=0.927 is a very close second. The fused images are shown in Figs. 7–9. Perceptually, the *a*DWTi fused images appear to look the best, although this was not measured empirically. Note that no postprocessing was imposed on these fused images.

Figure 5(a) shows the evaluation of function IQI (L_d,L_w) while fusing the medical image pair. Figure 5(b) shows the distributions of IQI (L_d) and IQI (L_w)—two slices of Fig. 5(a) at $L_w=4$ and $L_d=4$, respectively. Figure

(a) (b) (c) (d) (e)

Fig. 7 Image fusion with off-focus images (512×512 pixels): (a) input image A, (b) input image B, (c) image fused by iterative *a*DWTi, (d) image fused by Laplacian pyramid, and (e) image fused by regular DWT.

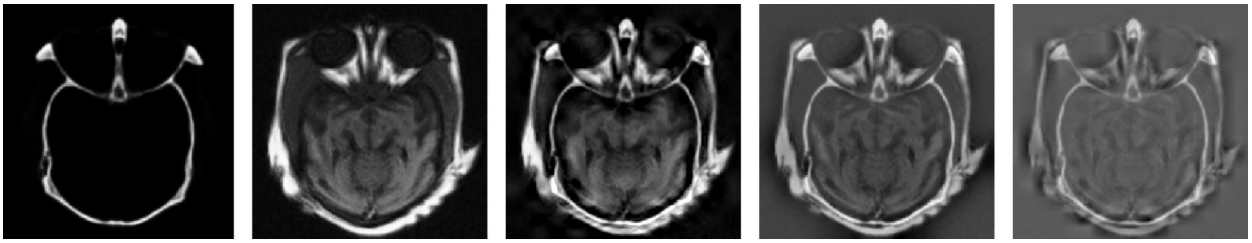

(a) (b) (c) (d) (e)

Fig. 8 Image fusion with medical images (256×256 pixels): (a) image A (CT), (b) image B (MRI), (c) image fused by iterative *a*DWTi, (d) image fused by Laplacian pyramid, and (e) image fused by regular DWT.

(a) (b) (c) (d) (e)

Fig. 9 Image fusion with remote sensing images (512×512 pixels): (a) image A (IR), (b) image B (low light), (c) image fused by iterative *a*DWTi, (d) image fused by Laplacian pyramid, and (e) image fused by regular DWT.

Table 3 The optimization traces of three metrics while fusing the medical pair with the iterative *a*DWT (*a*DWTi).

(L_d, L_w)	(1,4)	(2,4)	(3,4)	(4,4)	(5,4)	(6,4)	(5,2)	(5,6)	(5,8)	(**5,10**)
Entropy	6.171	6.419	6.607	6.787	*6.820*	6.825	6.834	6.844	6.850	**6.874**
SF	20.289	23.257	26.838	27.811	*28.227*	28.213	29.347	27.766	27.949	**27.918**
IQI	0.616	0.639	0.710	0.750	*0.772*	0.765	0.771	0.779	0.779	**0.780**

The *a*DWTi algorithm tried the fusion with parameters (6,4), but finally $L_d = 5$ was selected (with larger IQI) to enter next iterations of varying L_w. Only the evaluations of even numbers of L_w were given here to just simplify the table.

(a) (b) (c) (d) (e)

Fig. 10 Image fusion with the image pair NVESD10 (531×401 pixels) from NV-I1: (a) image A (II), (b) image B (IR), (c) image fused by iterative *a*DWTi (IQI=0.817), (d) image fused by Laplacian pyramid (IQI=0.716), and (e) image fused by regular DWT (IQI=0.565).

Table 4 The means and standard deviations (Std. dev.) of three metrics (entropy, SF and IQI).

Imagery	Statis.	Input A	Input B	Metric	Lapl.	Ratio	Contr.	Grad.	Morph.	DWT	*a*DWT	*a*DWTi
NV-I1 NVESD 24 pairs	Mean	5.968	6.833	Entropy	6.926	6.470	6.770	6.946	6.837	6.798	7.158	7.261
		10.277	13.750	SF	13.900	10.018	14.398	13.676	13.258	13.503	16.710	16.366
				IQI	0.772	0.495	0.672	0.722	0.739	0.640	*0.796*	**0.819**
	Std. dev.	0.625	0.237	Entropy	0.336	0.605	0.418	0.359	0.407	0.391	0.208	0.202
		2.853	1.780	SF	1.645	2.401	3.194	1.980	1.875	1.682	1.904	1.913
				IQI	0.062	0.112	0.101	0.065	0.064	0.060	0.079	0.070
NV-12 Lab 28 pairs	Mean	5.590	5.936	Entropy	7.001	6.512	6.636	7.075	6.857	6.898	7.077	7.177
		15.300	24.954	SF	20.999	15.209	22.022	20.570	20.483	19.387	28.166	27.773
				IQI	0.754	0.468	0.544	0.712	0.708	0.671	*0.766*	**0.788**
	Std. dev.	1.125	0.641	Entropy	0.458	0.784	0.858	0.432	0.432	0.417	0.584	0.538
		6.505	9.264	SF	6.888	4.334	7.221	6.597	6.079	6.012	9.345	9.064
				IQI	0.069	0.160	0.171	0.071	0.071	0.073	0.062	0.063

Eight fusion methods were tested on the NV-I1 (NVESD, 24 pairs) and the NV-I2 (Lab, 28 pairs). All seven noniterative fusion methods were done with default parameters $(L_d, L_w)=(4,4)$. The converged parameters of *a*DWTi (L_d, L_w) are dependent on the input image pair and the means of iterations are $N_i(\text{NV-I1})=13.833$, $N_i(\text{NV-I12})=13.607$.

5(b) shows an almost increasing function of IQI (L_d), while IQI (L_w) contains local fluctuations.

The optimization traces of the three metrics obtained after fusing the medical imagery pair with the iterative *a*DWT (*a*DWTi) algorithm are presented in Table 3. Fusion with the *a*DWTi algorithm was tried with parameters values of (6,4) during the iteration of varying L_d, but finally the best result (in the sense of IQI) given by $(L_d, L_w)=(5,4)$ was selected, the algorithm then entered the next iteration of varying L_w. Eventually, the maximum of IQI (=0.780) was chosen as the final result of the *a*DWTi.

4.1.3 *Two types of night vision imagery*

Two types of night vision imagery were tested in the current experiments: (1) 24 pairs of image intensified (II) and IR night-vision images (Army NVESD imagery, NV-I1) taken outdoors and (2) 28 pairs of near-IR and IR images (Lab imagery, NV-I2) that were obtained both indoors and outdoors (Indigo Systems, Inc., IR and near-IR Phoenix cameras). The source images were moderately noisy due to collection conditions. Therefore, to suppress the noise, a 2-D median filter (with a 3×3 sliding window) was first applied to all images before any fusion operation. This is a common preprocessing practice in fusing nighttime imagery; moreover, denoised images were used for all three algorithms (appropriate to compare the fused images). The filtered images are not shown here, but this filtering operation was not detrimental to the fused image quality [see Fig. 10(c) to (e)]. The complete assessments of the eight algorithms along with three evaluation measures are listed in Table 4. Here, the means and standard deviations of performance over all images in each group were computed. Although the image contents varied considerably, the averaged values still give a sense of an algorithm's performance. Keep in mind that the IQI values are always between [0,1] no matter what the analyzed image contents were. Judging from the IQI values, the *a*DWTi did the best, and the *a*DWT was the second best. Two typical pairs were drawn from each group and are illustrated in Figs. 10–13. The *a*DWTi results appear the best on causal visual inspection.

Since there were multiple image pairs for the night vision image sets, we were able to conduct a statistical analysis of the different IQI values produced by each fusion method examined in the current set of experiments. Specifically, a between-groups one-way analysis of variance (ANOVA) was performed on the IQI measures of both night vision imagery sets. Bonferroni-adjusted post-hocs were used for the multiple comparison portion of the analysis. The significant main effects were as follows: for NV-I1, $F(7,184)=42.327$, $p<0.0005$ and for NV-12, $F(7,216)=34.855$, $p<0.0005$. For NV-I1 (NVESD imagery), the performance of *a*DWTi was significantly better (p

Fig. 11 Image fusion with the image pair NVESD11 (531×401 pixels) from NV-I1: (a) image A (II), (b) image B (IR), (c) image fused by iterative *a*DWTi (IQI=0.795), (d) image fused by Laplacian pyramid (IQI=0.768), and (e) image fused by regular DWT (IQI=0.602).

Fig. 12 Image fusion with the image pair Lab09 (220×220 pixels) from NV-I2: (a) Image a (near IR), (b) image B (IR), (c) image fused by iterative *a*DWTi (IQI=0.756), (d) image fused by Laplacian pyramid (IQI=0.708), and (e) image fused by regular DWT (IQI=0.543).

<0.015) than that of other six methods except for *a*DWT. For NV-12 (Lab imagery), the performance of *a*DWTi was significantly better ($p<0.016$) than the other five methods except for *a*DWT and Laplacian pyramid.

4.2 *Discussion*

In the analysis described in the preceding sections, the metrics of entropy and SF were not described in detail. In fact, the SF is a more reliable and sensitive measure than the entropy. From the SF definition in Eq. (9), the value of SF reflects the detail or edge information contained in that image, thus it is sensitive to minor intensity changes in an image. Figure 14 shows the metric comparison between RMSE and SF, while fusing ''Lena-I1.'' The SF error (SFE) was obtained by subtracting SFs from the reference SF, 20.979. The smaller is the absolute SFE, the better is the fused image. It is clear that the SFE is very consistent with the RMSE. The problem is that there is no reference SF in a realistic image fusion application. On the other hand, the entropy is less sensitive to the detail changes of an image (refer to Table 1).

The iterative fusion algorithm, *a*DWTi, can achieve the optimal fusion (described by the IQI measurement, parameters, iterations, etc.) by adapting its parameter settings for the different input images. The *a*DWTi algorithm is obviously better than any of the noniterative algorithms examined here (including *a*DWT) judged on the IQI metric or visual inspection. In addition, the *a*DWTi is easily implemented to find the optimal IQI by separately varying two parameters (see Figs. 4 and 5), the DWT decomposition level (L_d) and the wavelet length (L_w), with no more than 15 iterations. And the number of iterations may be further reduced by properly initializing parameters or optimizing the termination conditions.

The evident drawback of using IQI as the feedback of an iterative fusion algorithm is that the error IQI ($=\mathrm{IQI}_F-1$) cannot provide the error direction (requirement 3 in Sec. 3.3) because the error IQI is always negative. As mentioned before, SF is a potential metric that can be developed to measure the error direction (see Fig. 14). We will report further work on how to use the SF measure as the feedback of an iterative fusion algorithm in a subsequent paper.

Overall, the iterative *a*DWT (*a*DWTi) produced the best results in fusing inhomogeneous imagery among others. Typically the noniterative *a*DWT was the second best. For off-focus images, both of which have similar brightness and contrast, the contrast pyramid could yield the best fusion among other seven methods. For night vision imagery where brightness and contrast were very different between the two input images, the Laplacian pyramid usually produced the third best fused image judged on the IQI metric, however, it often caused a visually overshot result in the high-contrast regions [see Fig. 10(d)].

4.3 *Future Work*

We plan to extend the definition of SF and develop its use as feedback of an iterative fusion algorithm. Importantly, we plan to conduct functional evaluation[5] to evaluate fusion performance based on human perceptual performance and compare the results to those from the quantitative metrics.

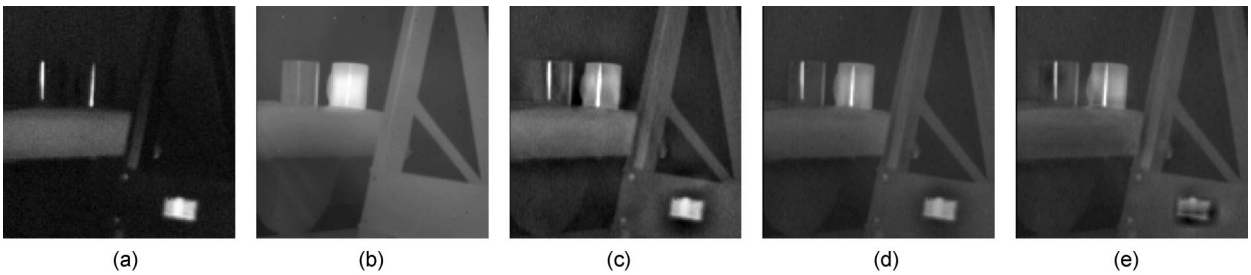

Fig. 13 Image fusion with the image pair Lab15 (220×220 pixels) from NV-I2: (a) image A (near IR), (b) image B (IR), (c) image fused by iterative *a*DWTi (IQI=0.852), (d) image fused by Laplacian pyramid (IQI=0.757), and (e) image fused by regular DWT (IQI=0.596).

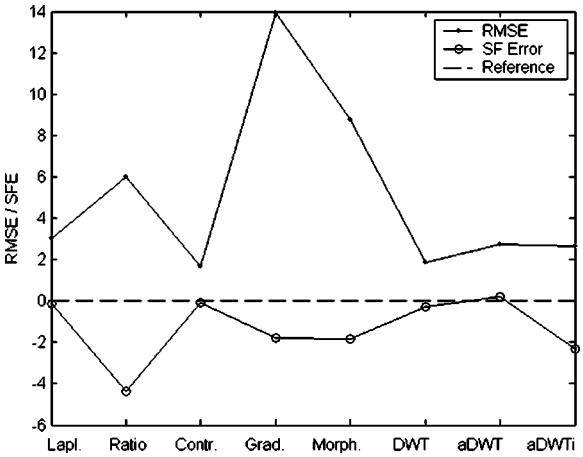

Fig. 14 Metric comparison between RMSE and SF while fusing "Lena-I1." The SF was obtained by subtracting SFs from the reference SF, 20.979. The smaller is the absolute SFE, the better is the fusion. It is clear that the SFE is very consistent with the RMSE.

5 Conclusions

We presented a new image fusion method, the advanced DWT (*a*DWT), which incorporated PCA and morphological processing into a regular DWT fusion procedure. Further, the *a*DWT was developed into an iterative fusion procedure directed by the metric of IQI. We compared the image fusion performance of six common methods (five pyramid methods and a regular DWT method) and our novel methods, the *a*DWT and the iterative *a*DWTi (*a*DWTi), based on four important quantitative measures—RMSE, the entropy, the SF, and the IQI. Overall, across the four different kinds of imagery, the *a*DWTi performed the best, and the *a*DWT was, for the most part, the second best compared the regular DWT or pyramid methods judged on the IQI metric and visual inspection of the fused imagery. Different image sources vary considerably in their intensities, contrast, noise, and intrinsic characteristics; therefore a large challenge for a fusion algorithm is to perform well across a variety of image sources, thus, *a*DWTi and *a*DWT are very promising to meet this goal.

Acknowledgment

This work was supported by Grant No. N00014-03-1-0224 from the Office of Naval Research.

References

1. J. L. VAN Genderen and C. Pohl, "Image fusion: Issues, techniques and applications," in *Intelligent Image Fusion, Proc. EARSeL Workshop*, J. L. van Genderen and V. Cappellini, Eds., pp. 18–26, Strasbourg, France (Sep. 1994).
2. L. D. Keys, N. J. Schmidt, and B. E. Phillips, "A prototype example of sensor fusion used for a siting analysis," in *Technical Papers 1990, ACSM-ASPRS Annual Conv. Image Processing and Remote Sensing* **4**, pp. 238–249 (1990).
3. R. H. Rogers and L. Wood, "The history and status of merging multiple sensor data: an overview," in *Technical Papers 1990, ACSM-ASPRS Annual Conv. Image Processing and Remote Sensing* **4**, pp. 352–360 (1990).
4. E. A. Essock, M. J. Sinai, J. S. McCarley, W. K. Krebs, and J. K. DeFord, "Perceptual ability with real-world nighttime scenes: image-intensified, infrared, and fused-color imagery," *Hum. Factors* **41**(3), 438–452 (1999).
5. E. A. Essock, J. S. McCarley, M. J. Sinai, and J. K. DeFord, "Human perception of sensor-fused imagery," in *Interpreting Remote Sensing Imagery: Human Factors*, R. R. Hoffman and A. B. Markman, Eds., Lewis Publishers, Boca Raton, Florida (Feb. 2001).
6. C. Pohl and J. L. Van Genderen, "Review article: multisensor image fusion in remote sensing: concepts, methods and applications," *Int. J. Remote Sens.* **19**(5), 823–854 (1998).
7. O. Rockinger, "Pixel level fusion of image sequences using wavelet frames," in *Proc. 16th Leeds Applied Shape Research Workshop*, pp. 149–154, Leeds University Press (1996).
8. H. Li, B. S. Manjunath, and S. K. Mitra, "Multisensor image fusion using the wavelet transform," *Graph. Models Image Process.* **57**(3), 235–245 (1995).
9. T. Pu and G. Ni, "Contrast-based image fusion using the discrete wavelet transform," *Opt. Eng.* **39**(8), 2075–2082 (2000).
10. D. A. Yocky, "Image merging and data fusion by means of the discrete two-dimensional wavelet transform," *J. Opt. Soc. Am. A* **12**(9), 1834–1841 (1995).
11. J. Nunez, X. Otazu, O. Fors, A. Prades, V. Pala, and R. Arbiol, "Image fusion with additive multiresolution wavelet decomposition; applications to spot+landsat images," *J. Opt. Soc. Am. A* **16**, 467–474 (1999).
12. F. Jahard, D. A. Fish, A. A. Rio, and C. P. Thompson, "Far/near infrared adapted pyramid-based fusion for automotive night vision," in *IEEE Proc. 6th Int. Conf. on Image Processing and its Applications (IPA97)*, pp. 886–890 (1997).
13. B. Ajazzi, L. Alparone, S. Baronti, and R. Carla, "Assessment of pyramid-based multisensor image data fusion," in *Proc. SPIE* **3500**, 237–248 (1998).
14. S. Li, J. T. Kwok, and Y. Wang, "Combination of images with diverse focuses using the spatial frequency," *Infusionstherapie* **2**(3), 169–176 (2001).
15. M. Beauchemin, K. B. Fung, and X. Geng, "A method based on local variance for quality assessment of multiresolution image fusion," *Proc. IAPRS*, Vol. XXXIV, p. B-32ff (2002).
16. L. W. Leung, B. King, and V. Vohora, "Comparison of image data fusion techniques using entropy and INI," in *Proc. ACRS*, Vol. 1, pp. 152–157 (2001).
17. G. Piella and H. Heijmans, "A new quality metric for image fusion," in *Proc. 2003 Int. Conf. on Image Processing*, Barcelona, Spain (2003).
18. Y. Zheng, E. A. Essock, and B. C. Hansen, "An advanced image fusion algorithm based on wavelet transform—incorporation with PCA and morphological processing," *Proc. SPIE* **5298**, 177–187 (2004).
19. P. J. Burt and E. H. Adelson, "Merging images through pattern decomposition," *Proc. SPIE* **575**, 173–182 (1985).
20. Z. Wang and A. C. Bovik, "A universal image quality index," *IEEE Signal Process. Lett.* **9**(3), 81–84 (2002).
21. A. M. Eskicioglu and P. S. Fisher, "Image quality measure and their performance," *IEEE Trans. Commun.* **43**(12), 2959–2965 (1995).
22. P. J. Burt, "The pyramid as structure for efficient computation," in *Multiresolution Image Processing and Analysis*, A. Rosenfeld, Ed., pp. 6–35, Springer-Verlag, New York/Berlin (1984).
23. P. J. Burt and E. Adelson, "The Laplacian pyramid as a compact image code," *IEEE Trans. Commun.* **Com-31**(4), 532–540 (1983).
24. A. Toet, "Image fusion by a ratio of low pass pyramid," *Pattern Recogn. Lett.* **9**(4), 245–253 (1989).
25. A. Toet, L. J. Van Ruyven, and J. M. Valeton, "Merging thermal and visual images by contrast pyramid," *Opt. Eng.* **28**(7), 789–792 (1989).
26. A. Toet, "Hierarchical image fusion," *Mach. Vision Appl.* **3**(1), 1–11 (1990).
27. A. Toet, "Multiscale contrast enhancement with application to image fusion," *Opt. Eng.* **31**(5), 1026–1031 (1992).
28. P. J. Burt, "A gradient pyramid basis for pattern-selective image fusion," *SID Int. Symp. Digest Tech. Papers* **16**, 467–470 (1985).
29. P. J. Burt and R. J. Lolczynski, "Enhanced image capture through fusion," in *IEEE Proc. 4th Int. Conf. Computer Vision*, pp. 173–182 (1993).
30. S. Richard, F. Sims, and M. A. Phillips, "Target signature consistency of image data fusion alternatives," *Opt. Eng.* **36**(3), 743–754 (1997).
31. A. Toet, "A morphological pyramid image decomposition," *Pattern Recogn. Lett.* **9**(4), 255–261 (1989).
32. G. K. Matsopoulos, S. Marshall, and J. Brunt, "Multiresolution morphological fusion of MR and CT images of the human brain," *IEE Proc. Vision Image Signal Process.* **141**(3), 137–142 (1994).
33. G. K. Matsopoulos and S. Marshall, "Application of morphological pyramids: fusion of MR and CT phantoms," *J. Visual Commun. Image Represent* **6**(2), 196–207 (1995).
34. R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, 2nd ed., Prentice Hall, Upper Saddle River, NJ (2002).

Yufeng Zheng received his PhD degree from the Tianjin University, China, in 1997. He is currently a postdoctoral research associate with the University of Louisville, Kentucky. His research interests are digital image processing, including image registration and fusion, noise reduction, image restoration and enhancement, edge detection, segmentation, and pattern recognition.

Edward A. Essock is a professor of Psychological and Brain Sciences and also of Ophthalmology and Visual Science with the University of Louisville. He is a vision scientist with broad research interests that include human factors and clinical applications of vision science, as well as basic science research. He received his training in vision science at Brown University, where he received his PhD degree in 1981, and as a postdoctoral fellowship at the University of California Berkeley with Jay Enoch in the School of Optometry. Dr. Essock has authored more than a 100 scientific papers, book chapters, and meeting publications. His current research focuses on detection and prediction of progression in glaucoma based on assessments of the nerve fiber layer thickness and biases in the perception of orientation in natural scenes (the ''horizontal effect'') and associated issues of cortical gain control.

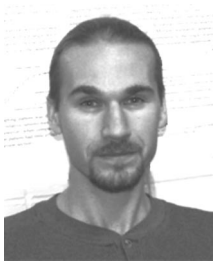

Bruce C. Hansen received his BS degree from the University of Michigan, Flint, in 2000, and his MA and PhD degrees from the University of Louisville, Kentucky, in 2003 and 2004, respectively, all in experimental psychology. He is currently a postdoctoral research associate with McGill University, Montreal, Quebec, Canada. His research interests include human visual processing of broadband spatial structure, multisensor image fusion, neural network modeling, multisensory integration, and human factors research.