

A region-based multiresolution image fusion algorithm*

Gemma Piella

Centre for Mathematics and Computer Science (CWI),
Amsterdam, The Netherlands
Gemma.Piella@cwi.nl

Abstract – *We propose a multiresolution fusion algorithm which combines aspects of region and pixel-based fusion. We use multiresolution decompositions to represent the input images at different scales, and introduce a multiresolution/multimodal segmentation to partition the image domain at these scales. The basic idea is to use this segmentation to guide the fusion process.*

A region-based multiresolution approach allows us to consider low-level as well as intermediate-level structures, and to impose data-dependent consistency constraints based on spatial, inter and intra-scale dependencies.

Keywords: image fusion, multiresolution decomposition, multimodal segmentation.

1 Introduction

Multiresolution (MR) techniques have proven to be very useful for image fusion. The algorithms based on MR techniques found in the literature are mainly pixel-based approaches where each individual coefficient of the MR decomposition (or possibly the coefficients in a small fixed window) is treated independently. However, in the majority of applications we are not interested in individual pixels but rather in the objects and regions they represent. It therefore seems reasonable to incorporate object and region information into the fusion process. Moreover, region-based approaches may help to circumvent some of the well-known drawbacks of pixel-based techniques, such as blurring effects, high sensitivity to noise and misregistration.

In this paper we propose a MR fusion algorithm which combines aspects of region and pixel-based fusion. The basic idea is to make a segmentation based on all different source images and to use this segmen-

tation to guide the fusion process. The paper is organized as follows. Section 2 briefly reviews some basic MR decomposition methods. Section 3 describes a general framework for pixel-based MR fusion. Extending this framework, Section 4 proposes a novel region-based MR fusion scheme. Some examples are also presented and discussed. Finally, Section 5 ends with some conclusions.

2 MR decomposition schemes

A MR scheme decomposes the signal being analyzed into several components, each of which captures information present at a given scale.

One of the earliest MR approaches is the pyramid decomposition scheme. A classical image pyramid [1] consists of a sequence of versions of an original image in which resolution is decreased in proportional steps. The bottom (or zero) level $x^{(0)}$ of the pyramid is equal to the original image x . This image is low-pass filtered and subsampled to obtain the next level $x^{(1)}$, which is then filtered and subsampled in the same way to obtain $x^{(2)}$. Further repetitions of this filtering/subsampling procedure generate the subsequent levels of the pyramid, also known as a *low-pass* or *approximation pyramid*, as each level contains successively coarser approximations of the original image. A difference pyramid can then be derived by interpolating each level of the low-pass pyramid and subtracting it from its predecessor. The resulting structure is called *high-pass* or *detail pyramid*, since each level contains only those features lost from one level to the next, and therefore contains only details within a restricted range of sizes.

Formally, the *pyramid transform* is the mapping $x^{(0)} \mapsto \{y^{(1)}, y^{(2)}, \dots, y^{(K)}, x^{(K)}\}$, where $y^{(k)}$ is the detail image at level k and $x^{(K)}$ is the approximation at the coarsest level K . Note that the pyramid transform of an image x constitutes a complete representation of x , since the steps used to construct the detail pyramid

*This paper has been published in ISIF Fusion 2002 conference, Annapolis, July 2002.

can be reversed to exactly recover x .

There are several modifications of the pyramid construction method described above. For instance, one can substitute the linear filters by morphological ones, which results in a morphological pyramid [2]; alternatively one can compute the ratio (instead of the difference) of the successive low-pass filtered images to construct the so-called ratio of low-pass pyramid [3]. Another alternative is the gradient pyramid [4], which results when the gradient operator is applied to each level of the low-pass pyramid. In this case, the detail images $y^{(k)}$ are further separated into subbands according to their orientation; namely, $y^{(k)} = \{y^{(k)}(\cdot|1), y^{(k)}(\cdot|2), y^{(k)}(\cdot|3), y^{(k)}(\cdot|4)\}$, representing the horizontal, vertical and the two diagonals directions.

Another MR analysis method similar to pyramids is the *wavelet decomposition* [5]. The main difference is that while pyramids lead to an overcomplete set of transform coefficients, the wavelet decomposition results in a non-redundant representation, i.e., the size of the input equals the size of the output.

Most of the MR transforms used in practice are one-dimensional. In many cases, one uses separable wavelet decompositions for images. That means that an image is first decomposed row-wise by a one-dimensional transform; the resulting image is then decomposed column-wise by the same (or another) one-dimensional transform. This results in four images at the first level, one approximation image and three detail images. This procedure can be repeated on the approximation image. At each level k , the input $x^{(k)}$ is decomposed into a coarser approximation $x^{(k+1)}$ and three detail signals $y^{(k+1)} = \{y^{(k+1)}(\cdot|1), y^{(k+1)}(\cdot|2), y^{(k+1)}(\cdot|3)\}$, corresponding to the horizontal, vertical and diagonal directions.

One drawback of the (discrete) wavelet transform is that it yields a shift variant signal representation. This means that a simple shift of the input signal may lead to completely different transform coefficients. The lack of translation invariance can be avoided if the outputs of the filter banks are not decimated. The undecimated wavelet transform [5] then produces a set of band-pass signals which have the same size as the original input.

Non-separable transforms can also be constructed and provide shift invariance and/or directional selectivity. Moreover, nonlinear extensions of the wavelet transform are possible [6, 7]. A general and flexible approach for the construction of MR decomposition schemes can be found in [7, 8].

In this paper, the MR decomposition of an image $x = x^{(0)}$ is denoted by y and it is assumed to be of the

form:

$$y = \{y^{(1)}, y^{(2)}, \dots, y^{(K)}, x^{(K)}\}. \quad (1)$$

Here $x^{(K)}$ represents the approximation image at the highest level (lowest resolution) of the MR structure, while images $y^{(k)}$, $k = 1, \dots, K$, represent the detail images at level k . The detail at level k will, in general, comprise various frequency or orientation bands, depending on the type of MR transform that has been used. We assume henceforth that $y^{(k)}$ is composed of P detail images: $y^{(k)} = \{y^{(k)}(\cdot|1), \dots, y^{(k)}(\cdot|P)\}$. We use the coordinates (i, j) to index the spatial position within each level k and frequency band p . Thus, $x^{(k)}(i, j)$ represents the approximation coefficient at spatial position (i, j) within the level k . Similarly, $y^{(k)}(i, j|p)$ represents the detail coefficient at spatial position (i, j) within the level k and the band p .

3 Generic pixel-based MR fusion scheme

The general structure of fusion schemes based on MR decompositions is depicted in Fig. 1. First, a MR transform is applied to all source images. Then, a composite MR representation is constructed by a combination of the MR coefficients of the sources. Finally, the fused image is obtained by computing the inverse transform of the composite representation.




Figure 1: *MR image fusion methodology: (i) MR transform (Ψ) of the sources, (ii) combination in the transform domain and (iii) inverse MR transform (Ψ^{-1}) of the combined representation.*

Several implementations of this general scheme exist in the literature [9, 10, 11, 12, 13]. They differ in the type of MR transform employed and in the combination algorithm. A general framework which encompasses most of the existing MR-based fusion schemes is shown in Fig. 2. We describe each of its parts in some detail.

MR analysis (Ψ)

As the name suggests, this block performs a MR decomposition of the input. Some issues to be addressed at this stage are the specific type of decomposition and the number of levels.

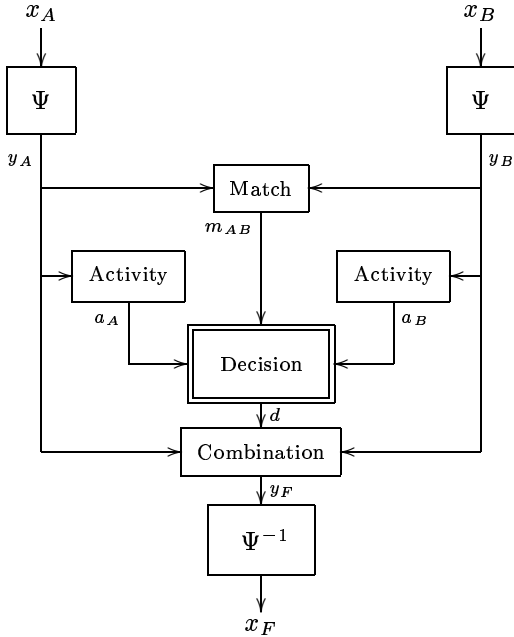


Figure 2: *Generic pixel-based MR fusion scheme with two input sources x_A and x_B , and one output fused image x_F .*

We assume that the same type of transform Ψ and number of levels K are applied to all sources x_S , $S \in \mathcal{S}$, where \mathcal{S} is the index set of source images. Thus, for every input x_S we obtain its MR representation $y_S = \Psi(x_S)$, with y_S having the form defined in (1). For convenience, we will denote the approximation image $x_S^{(K)}$ by $y_S^{(K)}(\cdot|0)$. In this way, we can use the general expression of $y_S^{(k)}(\cdot|p)$ to refer both to the detail images (for $p = 1, \dots, P$) and the approximation image (for $p = 0$).

As we have seen in the previous section, the MR representations y_S comprises information at different scales. High levels contain coarse scale information while low levels contain finer details. Such representation is suitable for image fusion, not only because it enables one to consider and fuse image features separately at different scales, but also because it produces large coefficients near edges, thus revealing salient information.

Activity

This block associates to every image $y_S^{(k)}(\cdot|p)$ an activity level $a_S^{(k)}(\cdot|p)$, which reflects the local activity of the image. Broadly speaking, the activity level $a_S^{(k)}(i, j|p)$ of a sample (i, j) will be high if the average energy (or some other measure) of $y_S^{(k)}(\cdot|p)$ is high in the vicinity of (i, j) . Thus, one may assume that

the activity level depends on some sort of energy calculation over a local neighbourhood of coefficients. In the simplest case, the activity level is just the absolute value of the coefficient, that is,

$$a_S^{(k)}(i, j|p) = |y_S^{(k)}(i, j|p)|. \quad (2)$$

Match

This measure is supposed to quantify the degree of ‘similarity’ between the sources. More precisely, the match value $m_{AB}^{(k)}(i, j|p)$ reflects the resemblance between the inputs $y_A^{(k)}(i, j|p)$ and $y_B^{(k)}(i, j|p)$. The relative amplitudes of the samples can be used as a measure of their similarity or match. Alternatively some other criteria can be used. In the following expression, the match value is a normalized correlation averaged over a neighbourhood of the samples:

$$m_{AB}^{(k)}(i, j|p) = \frac{2 \sum_{(u,v) \in \mathcal{W}} y_A^{(k)}(u, v|p) y_B^{(k)}(u, v|p)}{|y_A^{(k)}(i, j|p)|^2 + |y_B^{(k)}(i, j|p)|^2}, \quad (3)$$

where \mathcal{W} is a window centered at (i, j) .

By analyzing the match value, one can determine where the sources differ and to which extent, and use this information to combine them in an appropriate way. For example, if the match measure at a given position is low (i.e., the sources are distinctly different at that position), the coefficient from the source decomposition with the highest activity level is taken as the composite coefficient. On the other hand, if the match value is high (i.e., the sources are similar at that position), the coefficients from the different sources are averaged to yield the composite coefficient.

Decision

This block is the core of the combination algorithm. It decides about the actual combination of the coefficients of the MR decompositions of the various sources. The activity level of every source together with the match measure are used as input for the decision process.

For each level k , orientation band p and sample position (i, j) , the decision process assigns a decision parameter $\delta = d^{(k)}(i, j|p)$ which will determine how to compute the composite $y_F^{(k)}(i, j|p)$. Most often, the decision takes place independently at each level, band and position. However, it may also take into account spatial, inter and intra-scale dependencies between the samples, thus exploiting the idea that coefficients in the composite should not be computed independently. For instance, one may require neighbouring coefficients in the same level and/or orientation to take the same decision.

Combination

This module describes the actual combination of the transform coefficients of the sources. For simplicity, let us consider two sources and let us assume that every composite coefficient $y_F^{(k)}(i, j|p)$ is ‘assembled’ from the source coefficients at the corresponding level, band and position. More precisely,

$$y_F^{(k)}(i, j|p) = C^{(k)}\left(y_A^{(k)}(i, j|p), y_B^{(k)}(i, j|p), \delta\right), \quad (4)$$

where $C^{(k)} : \mathbb{R}^3 \mapsto \mathbb{R}$ is the combination function parametrized by the decision $\delta = d^{(k)}(i, j|p)$. A simple choice for $C^{(k)}$ is a linear mapping, e.g.,

$$C^{(k)}(a, b, \delta) = w_A(\delta)a + w_B(\delta)b, \quad (5)$$

where the weights $w_A(\delta)$, $w_B(\delta)$ depend on the decision δ .

MR synthesis (Ψ^{-1})

Finally, the fused image is obtained by applying the inverse transformation on the composite MR decomposition y_F :

$$x_F = \Psi^{-1}(y_F), \quad (6)$$

where Ψ^{-1} is the inverse MR transform.

Let us assume that the composite coefficients $y_F^{(k)}(i, j|p)$ are obtained by a linear combination as in (5). It remains to be specified how the decision $\delta = d^{(k)}(i, j|p)$ determines the weights $w_S(\delta)$. A natural approach is to assign to each coefficient a weight that depends increasingly on the activity level. In general, the resulting weighted average leads to a stabilization of the fusion result, but it introduces the problem of contrast reduction in case of opposite contrast in different source images. This can be avoided by using a selective combination rule where the most salient component, i.e., the one with largest activity level, is chosen for the composite. In this case,

$$y_F^{(k)}(i, j|p) = y_M^{(k)}(i, j|p)$$

with $M = \arg \max_{S \in \mathcal{S}} \left(a_S^{(k)}(i, j|p)\right)$.

In many approaches, the composite approximation coefficients of the highest decomposition level, representing the mean intensity, are taken to be a weighted average of the approximation of the sources, while the composite detail coefficients use a selective combination. Additionally, a match measure can be used to determine which combination mode, averaging or selection, to use. Moreover, the decision can be made globally for a group of samples; for example, constraining the samples in the same level k and position (i, j) to take the same decision for all bands p , or ‘extending’ a decision to a spatial neighbourhood.

Note finally that other factors may be incorporated for the fusion rules. In particular, if some prior knowledge is available, all the fusion blocks can use such information to improve fusion performance. For instance, when combining the source coefficients, the weights assigned to them may depend not only on the activity level and match measure, but may also reflect some a-priori knowledge of a specific type, giving preference to certain levels k , spatial positions (i, j) or some input sources.

4 Region-based MR fusion

4.1 Motivation

As we pointed out before, many of the existing MR fusion schemes fit in the general scheme depicted in Fig. 2 and described in Section 3. These schemes are pixel-based in the sense that each transform coefficient (or possibly the coefficients within a small fixed neighbourhood) is considered separately.

For most, if not all, image fusion applications, it could be more meaningful to combine objects rather than pixels. As an intermediate step from pixel-based towards object-based fusion schemes, one might consider region-based approaches. This has the additional advantage that the fusion process becomes more robust and avoids some of the problems in pixel-level fusion, such as high sensitivity to noise and blurring effects.

4.2 The general scheme

Our region-based fusion scheme (see Fig. 3) extends the pixel-based fusion approach discussed in Section 3. Indeed, it includes all the blocks described before. In the present scheme the input sources x_S are also used to compute a single MR segmentation \mathcal{R} . That is, we use MR decompositions to represent the input images at different scales and, additionally, we introduce a multiresolution/multimodal (MR/MM) segmentation to partition the image domain at these scales. The activity level and match measures are computed for every region in the decomposed input images. These measures may correspond to low-level as well as intermediate-level structures. Furthermore, the MR segmentation \mathcal{R} allows us to impose data-dependent consistency constraints based on spatial as well as inter and intra-scale dependencies. All this information, i.e, the measures and the consistency constraints, is integrated to yield a decision map d which governs the combination of the coefficients of the transformed sources. This combination results in a MR decomposition y_F , and by MR synthesis we obtain a fused image x_F .

The main functional blocks of this fusion strategy are depicted in Fig. 3. Since we already discussed most of

them in Section 3, we concentrate on the segmentation module and its interaction with the other modules.

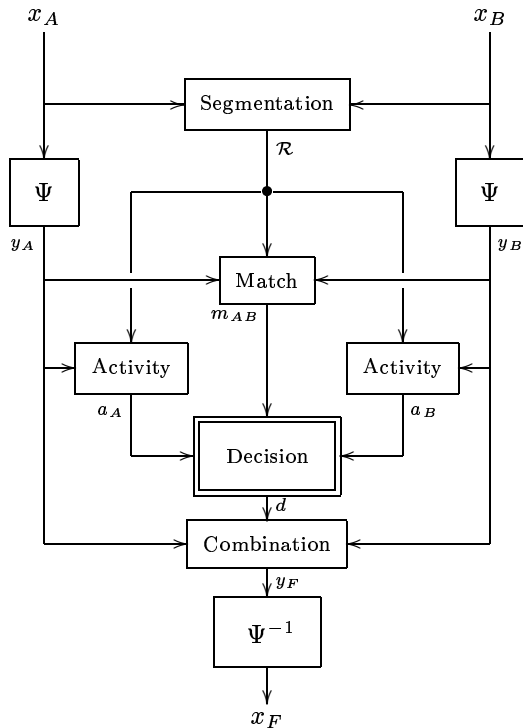


Figure 3: *Generic region-based MR fusion scheme with two input sources x_A and x_B , and one output fused image x_F .*

Segmentation

This module uses the various source images as input and returns a single MR segmentation

$$\mathcal{R} = \{\mathcal{R}^{(1)}, \mathcal{R}^{(2)}, \dots, \mathcal{R}^{(K)}\}$$

as output. Here $\mathcal{R}^{(k)}$ represents a segmentation at level k . The basic idea underlying such a segmentation is to identify the different regions in the scene in order to make a global interpretation of the data at a more abstract level. For this purpose, we develop a MR/MM segmentation algorithm [14] based on a generalized pyramid linking method [15]. Basically, the pyramid linking segmentation algorithm consists of four steps:

1. *Initialization.* Images at increasingly larger scales are obtained by a filtering/subsampling procedure to construct an approximation pyramid as described in Section 2. We impose the requirement that the sampling scheme used in this stage is the same as the one used in the MR analysis block. However, we emphasize that the approximation

pyramid constructed here may be completely different from the MR decomposition obtained in the MR analysis module.

2. *Linking.* Child-parent relationships are established between samples at adjacent resolution levels. For each child a suitable parent is sought among the candidate parents: it is linked to its most ‘similar’ parent. We define this similarity based on geometrical and grayscale proximity.
3. *Root labeling.* Samples having ‘weak’ parent links and all samples in the highest level are marked as roots.
4. *Downward projection.* The original image is segmented by tracing back the linkages from the roots to the bottom level. In each level, all the samples that are connected to a common root are classified as a single region segment R . Thus, at each level k , we obtain a segmented image $\mathcal{R}^{(k)}$ which contains the different regions R at this level.

Loosely speaking, \mathcal{R} provides a MR representation of the various regions of the underlying scene. This representation will guide the other blocks of the fusion process. In principle, these blocks are the same as in the pixel-based case. However, instead of working at pixel-level, we take into consideration the regions inferred by the segmentation. In the simplest case, these regions can be considered as elementary units and thus, instead of fusing pixel by pixel, we would be fusing region by region.

As in the pixel-based scheme, once the decision map is constructed, the mapping performed by the combination process is determined for all coefficients, and the synthesis process yields the fused image x_F .

Note that the fusion scheme (for the pixel as well as for the region-based case) encompasses several blocks, some of which are optional. Furthermore, the parameters and functions comprised by the various blocks can be chosen in a variety of ways. Since different combinations will lead to different performances, it is important to study and understand the effect of these choices on the final fusion process.

4.3 Case studies

Our current implementation of the proposed algorithm is still in an early stage of development. In this section, we consider one of the simplest approaches. Take two input images x_A and x_B . For their MR decomposition, we use a Laplacian pyramid (thus, we only have a single orientation band, i.e., $P = 1$, and we will omit the index p unless it explicitly equals 0). In the MR/MM segmentation algorithm, both the linking and root labeling criteria are based only on

the difference in grayscale values between parent and child. More sophisticated criteria, as well as connectivity or entropy considerations, are still to be incorporated. The combination process is performed as in (5), with $w_A(\delta) = \delta$ and $w_B(\delta) = 1 - \delta$. In the decision process, each component of d is obtained by the following simple decision rules:

- For $p = 0$,

$$\delta = d^{(K)}(i, j|0) = \frac{1}{2}, \text{ for all } (i, j).$$

- For $p = 1$,
- for each level k and for each region $R \in \mathcal{R}^{(k)}$

1. Compute for each y_S the activity level:

$$a_S^{(k)}(R) = \frac{1}{|R|} \sum_{(i,j) \in R} |y_S^{(k)}(i, j)|,$$

where $|R|$ is the area of region R .

2. Compute decision

$$d^{(k)}(R) = \begin{cases} 1 & \text{if } a_A^{(k)}(R) > a_B^{(k)}(R) \\ 0 & \text{otherwise.} \end{cases}$$

Thus, according to our algorithm, the composite approximation image $y_F^{(K)}(\cdot|0)$ is constructed by the (pixelwise) average of the approximation images $y_S^{(K)}(\cdot|0)$. In this case, the region information is neglected. The composite detail images $y_F^{(k)}$, however, are constructed by a selective combination where the composite coefficients of all samples belonging to a certain region come from the same source detail image (either $y_A^{(k)}$ or $y_B^{(k)}$), namely, the one whose activity in that region is higher. Obviously, this is a rather naive approach and we are currently exploring various alternatives.

We have tested our algorithm on several pairs of images. Three examples are given here to illustrate the fusion process described above. In all cases, we have chosen $K = 3$ and, when displaying the images, the gray values of the pixels have been scaled between 0 and 255 (histogram stretching). The input sources x_A and x_B are displayed, respectively, on the left and right top of the corresponding figure. For the decision maps, pixels with $\delta = 0, 1$ are displayed in black and white, respectively. Thus, according to our algorithm, coefficients corresponding to ‘white zones’ are selected from $y_A^{(k)}$, while coefficients corresponding to ‘black zones’ are selected from $y_B^{(k)}$.

Fig. 4 shows the fusion of images obtained with different sensors. Note that in the visual image (Fig. 4, top left) it is hard to distinguish the person in camouflage from the background, while it is clearly observable

in the infrared (IR) image (Fig. 4, top right). In contrast, the easily discernable background in the visual image is nearly undetectable in the IR image.




Figure 4: *Example 1. Top: visual (left) and IR (right) test images; middle: 1st level of segmentation (left) and decision (right); bottom: 2nd level of segmentation (left) and decision (central), and fused image (right). Source images are courtesy of Alex Toet, from TNO Human Factors Institute, The Netherlands.*

The first level of the resulting segmentation and decision map are shown in the middle row. The corresponding second levels are displayed on the bottom left. It is interesting to note that, according to $d^{(2)}$, although most of the background is selected from the visual image $y_A^{(2)}$, the region corresponding to the person is selected from the IR image $y_B^{(2)}$. The final fused image is depicted at the bottom right of Fig. 4.

The next example illustrates the usage of fusion in radiotherapy and skull surgery. Here, the information provided by magnetic resonance imaging (MRI) and X-ray computed tomography (CT) is complementary. Normal and pathological soft tissues are better visualized by MRI (Fig. 5, top left), while the structure of the tissue bone is better visualized by CT (Fig. 5, top right). The fused image, depicted at the bottom right of Fig. 5, provides salient information from both images simultaneously, which may be useful for physicians in medical diagnosis. As in the previous example, the segmentation and decision map for the first and

second level are also displayed.




Figure 5: *Example 2. Top: MRI (left) and CT (right) test images; middle: 1st level of segmentation (left) and decision (right); bottom: 2nd level of segmentation (left) and decision (center), and fused image (right).*

Fig. 6 shows how fusion can be used to extend the effective depth-of-field of a vision system. Due to the limited depth-of-field of optical lenses, it is often not possible to get an image with all objects in focus. One way to overcome this problem is to take several recordings with different focus points and combine them into a single composite which contains the focused regions of all input images. This could be useful, for example, in digital camera design or in industrial inspection applications where the necessity to visualize objects at very short distances complicates the preservation of the depth-of-field. The segmentation and decision for the first level are displayed in the second row of Fig. 6. Note that since the digit ‘8’ is connected to a particular region located within the left clock, the binary decision maps points out, wrongly, to take the ‘8’ from $y_A^{(1)}$ instead from $y_B^{(1)}$. For this particular example,

we also show the corresponding fused output we would have obtained using a pixel-based MR fusion algorithm with the same fusion rules as in the region-based algorithm. We also show how we can improve the region-based fused image by filtering the decision map. Here, we have used a morphological alternating filter: an opening (to remove small ‘white regions’) followed by a closing (to remove small ‘black regions’). The resulting filtered decision map and the corresponding fused image are displayed at the bottom of Fig. 6.

4.4 Discussion

Despite the crudeness of the current implementation, the visual performance is surprisingly good. This suggests that the region-based approach proposed here can at least be competitive with (but more likely outperform) other MR fusion techniques.

Further investigations are necessary for the fine-tuning of parameters as well as the proper selection of the different ingredients of the scheme. Towards this end, performance assessment criteria will be developed to evaluate and demonstrate the capacities of the new fusion technique, as well as to compare its performance with others MR fusion schemes. Test will be carried out based both on objective and subjective criteria.

5 Conclusions

In this paper, we have introduced a general framework for pixel-based MR image fusion. Extending this framework, we have proposed a region-based MR fusion scheme which combines aspects of both object and pixel-level fusion.

The implementation of our algorithm is still in a preliminary stage and in the experiments performed we did not attempt to optimize its performance. However, the results obtained so far suggest that our approach may be useful for several image fusion applications. We intend to investigate this more thoroughly in the future. A substantial part of our efforts will be devoted to the design of objective measures for fusion performance assessment.

Acknowledgments

This work is supported by the Dutch Technology Foundation STW, project no. CWI 4616. The author thanks A. G. Steenbeek and P. M. de Zeeuw for their help in the software implementation, and to A. Toet for providing some of the test images. The advice and support of Henk Heijmans is also gratefully acknowledged.




Figure 6: *Example 3. Top: multi-focus test images; second row: 1st level of segmentation (left) and decision (right); third row: fused images with pixel-based (left) and region-based (right) approach; bottom: filtered decision (left) and corresponding fused image.*

References

- [1] P. J. Burt and E. H. Adelson, "The Laplacian pyramid as a compact image code," *IEEE Transactions on Communications*, vol. 31, pp. 532–540, 1983.
- [2] A. Toet, "A morphological pyramidal image decomposition," *Pattern Recognition Letters*, vol. 9, pp. 255–261, 1989.

- [3] A. Toet, "Image fusion by a ratio of low-pass pyramid," *Pattern Recognition*, vol. 9, pp. 245–253, 1989.
- [4] P. J. Burt, "A gradient pyramid basis for pattern selective image fusion," in *Proceedings of the Society for Information Display Conference*, 1992.
- [5] S. G. Mallat, *A Wavelet Tour of Signal Processing*, Academic Press, San Diego, California, 1998.
- [6] R. L. de Queiroz, D. A. F. Florêncio, and R. W. Schafer, "Nonexpansive pyramid for image coding using a nonlinear filterbank," *IEEE Transactions on Image Processing*, vol. 7, pp. 246–252, 1998.
- [7] H. J. A. M. Heijmans and J. Goutsias, "Nonlinear multiresolution signal decomposition schemes. Part II: morphological wavelets," *IEEE Transactions on Image Processing*, vol. 9, no. 11, pp. 1897–1913, 2000.
- [8] J. Goutsias and Henk J. A. M. Heijmans, "Multiresolution signal decomposition schemes. Part I: Morphological pyramids," *IEEE Transactions on Image Processing*, vol. 9, no. 11, 2000.
- [9] P. J. Burt and R. J. Kolczynski, "Enhanced image capture through fusion," in *Proceedings of the 4th International Conference on Computer Vision*, Berlin, Germany, May 1993, pp. 173–182.
- [10] H. Li, B. S. Manjunath, and S. K. Mitra, "Multisensor image fusion using the wavelet transform," *Graphical Models and Image Processing*, vol. 57, no. 3, pp. 235–245, May 1995.
- [11] I. Koren, A. Laine, and F. Taylor, "Image fusion using steerable dyadic wavelet transforms," in *Proceedings of the IEEE International Conference on Image Processing*, Washington D.C., October 1995, pp. 232–235.
- [12] T. A. Wilson, S. K. Rogers, and L. R. Meyers, "Perceptual based hyperspectral image fusion using multiresolution analysis," *Optical Engineering*, vol. 34, no. 11, pp. 3154–3164, 1995.
- [13] L. J. Chipman and T. M. Orr, "Wavelets and image fusion," in *Proceedings of the IEEE International Conference on Image Processing*, Washington D.C., October 1995, pp. 248–251.
- [14] G. Piella, "Multiresolution image fusion: a general mathematical framework encompassing a new region-based approach," Research report, CWI, Amsterdam, 2002, in preparation.
- [15] P. J. Burt, T. H. Hong, and A. Rosenfeld, "Segmentation and estimation of image region properties through cooperative hierarchical computation," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 11, no. 12, pp. 802–809, December 1981.