A Hierarchical Approach for Registration of High-Resolution Polarimetric SAR images

Dirk Borghys and Christiaan Perneel and Marc Acheroy

Royal Military Academy, Signal & Image Centre, Renaissancelaan 30, B-1000 Brussels, Belgium

ABSTRACT

The presented work aims to automatically register high-resolution polarimetric SAR images with each other and other types of images. A digital topographic map is used as an aid for the registration. SAR images are very different from visual or infrared images. The idea is to identify, for each type of image, objects present on the map and easily detectable in the image. Detecting these objects in the image and matching them between map and image provides a first registration. Several object detectors were developed for the subsequent stages of the registration. Each of these detectors is briefly described. The actual registration uses a hierarchical method. First the SAR image is converted into ground range. Then a rough registration between image and map is obtained based on the position of forests and/or built-up areas. A voting method is used to find the parameters of a simple transformation model and to match the objects between map and image. The third step finds the parameters of an affine transformation based on the objects matched by the voting method. To improve the registration, objects with low 3D structure, e.g. roads and rivers, are used. The method for detecting these in SAR images yields an incomplete results leading to ambiguities for the optimal local displacement. Optimisation methods are used to overcome this problem and yield the parameters of a global transformation model. The accuracy of the registration is now within the accuracy of the map. Once the different images are registered with the map, the results of edge detectors are used to refine the registration between them.

Keywords: Image Registration, Polarimetric SAR, Built-up Areas, Edge Detection

1. INTRODUCTION

Results of remote sensing applications can be greatly improved when multi-sensor data are used. In order to combine the information contained in the different images, it is necessary to register them. As the range of different sensor types and the spatial resolution of the images increase, automatic registration becomes more difficult because it needs to be able to cope with very different images and the registration accuracy must be very high. This is particularly true when trying to register SAR images with optical or infrared images. A method to register SAR images based on patch (uniform region) matching between different images is proposed in.[1,2] The method starts from a initial, rough manual registration indicating 4 control points on each image. Although the flight parameters of the sensor platform could allow to find the equivalent of such a manual registration, we have experienced that these flight parameters can be very inaccurate (or sometimes unavailable).

This paper proposes a registration method that only supposes that the spatial resolution of the image is roughly known and that the location in the scene is known up-to an accuracy of about 1 km. The proposed method uses a digitised topographic map as a guide in the registration. The idea is to identify, for each type of image, objects that are present on the map and easily detectable in the respective image. This leads to a first registration between each image and the map. The image-to-image registration is then a refinement based on features that are present on the different images. The method is applied to the registration of two high-resolution polarimetric SAR images. Fig. 1 presents the method used to register different SAR images. Each SAR image is first transformed to ground range and rescaled to approximately square pixels. The registration with the map consists in two steps. The first step uses the position of built-up areas and/or forests. Such objects have the advantage of being usually discriminative and well distributed over a given region. However, they are difficult to localise precisely in a SAR image due to their 3D structure. The second step therefore uses communication lines (roads, railways and waterways) to refine the registration. For the registration with the map an affine

Further author information: Send correspondence to Dirk.Borghys@elec.rma.ac.be

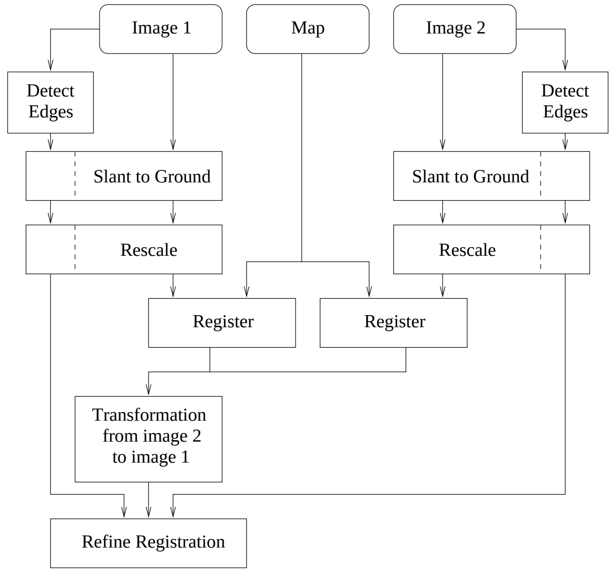

Figure 1. Overview of the registration method for SAR images

transform is used. The registration transformations from each image to the map can thus be easily inverted and combined. This allows to find an affine transformation that warps one image onto the other. The image-to-image registration is then just a refinement of this result. It is based on the position of edges in the different images. Although the method is applied to register two SAR images, it should be applicable to other types of images. The only components that need to be adapted are the algorithms used to detect the various objects in the image.

In section 4 the registration between image-to-map registration method is explained. Section 5 the image-to-image registration is described. In section 2 the various feature and object detectors for SAR image are briefly described and references to a more complete description are given. Section 3 proposes a quantitative evaluation method for image registration.

2. OVERVIEW OF THE FEATURE/OBJECT DETECTORS

The objects and features that are used in the registration method are built-up areas, forests, edges and communication lines. Specific detectors were developed to detect each of these in high-resolution polarimetric SAR images. For detecting built-up areas a feature based method was used.[3,4] Different statistical and polarimetric features were combined using logistic regression. The detection of forests is based on polarimetric decomposition methods. In[5] a review of such methods is presented. We have used the fuzzy-logic based classification based on decomposition methods proposed by M. Hellmann.[6] The edge detectors are based on statistical hypothesis tests.[7,8] The communication lines were detected by constructing a bar detector from the developed edge detectors in a similar way as in.[9]

3. EVALUATION METHOD FOR THE REGISTRATION

For the evaluation of the registration a set of corresponding control points is manually selected on the images to be registered (or the image and the map). The transformation that is found by the automatic registration is then used to warp the control points from the first image into the second and the distance between these warped points and the control points in the second image is used to define the evaluation criterion for the quality of registration. Let $P^m(i), i = 1..N$ be the control points in the first image and $P^s(i), i = 1..N$ the corresponding points in the second image. Φ is the transformation resulting from the automatic registration. For each point in the set of control points the distance is then calculated between the $P^s(i)$ and the warped point $\Phi(P^m(i))$:

$$\forall i = 1..N \quad d_i = Dist(P^s(i), \Phi(P^m(i))) \tag{1}$$

The three measurements used to estimate the quality of the registration are the maximum, the average and the root-mean-square of this distance:

$$D_{Max} = \max_{i=1..N} d_i, \qquad D_{Mean} = \frac{1}{N}\sum_{i=1}^{N} d_i, \qquad D_{RMS} = \frac{1}{N}\sqrt{\sum_{i=1}^{N} d_i^2} \tag{2}$$

4. REGISTRATION BETWEEN IMAGE AND MAP

The global strategy used for registering SAR images with a map is schematically presented in fig. 2. The refinement using contours of forests is optional. The conversion from slant range to ground range and the resampling to approximately square pixels is performed in a pre-processing step that relies only on a priori information about the image acquisition. This pre-processing is explained in the next section. The following sections respectively present the rough registration based on the feature consensus method and the refinements of the results using contours of forests and the network of communication lines.

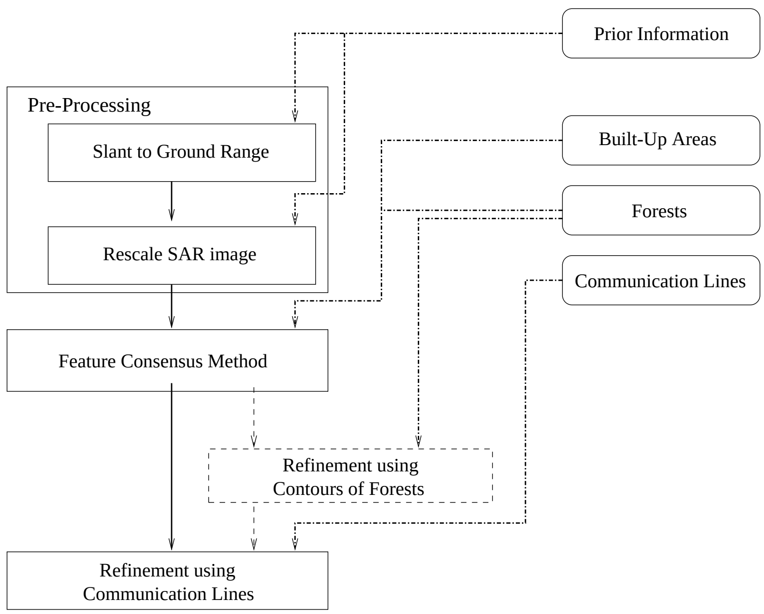

Figure 2. Registration of SAR Images with a Map: Overview of the Method

4.1. Pre-Processing based on A Priori Information

The pre-processing includes the conversion of the SAR image from slant range to ground range and its resampling to (almost) square pixels. Both steps are solely based on prior information about the image acquisition. They do not involve any matching process or an iterative search for some optimal transformation. This is the reason that we call it pre-processing. The conversion to ground range is based on the assumption of a flat terrain. At this point of the registration process it is not possible to incorporate any available terrain elevation data because we have not yet geocoded the image. In order to facilitate the feature consensus method, the SAR image is also resampled such that the pixel spacing is similar to the pixel spacing in the map.

4.2. Registration by Feature Consensus

The method used to find a first registration based on the position of built-up areas and forests is based on a technique called feature consensus that was first described in an article by Shekhar, Govindu and Chellappa[10] and refined by the same authors in.[11,12] Their approach does not require feature correspondence or area correlation. The geometric transformation to be used is reparametrised into a sequence of elementary stages. At each stage a single transformation parameter is estimated from the available features by a method called feature consensus. The method consists in choosing the parameter that is most consistent with all possible pairs of features, i.e. for each feature pair, the transformation that would map them onto each other is determined. This means that the parameter must be "visible" from the features that are used, e.g. when finding a rotation angle, it is not possible to use single points as features; one should at least use line segments. For all true feature pairs, the estimated parameters should cluster around the true value of the parameter while for all other pairs, parameters will be randomly distributed. The parameter is then chosen by selecting the maximum of the distribution of the parameters. The distribution for a given parameter is called the *consensus function* for that parameter. The method allows to find a global transformation only if it can be reparametrised as a sequence of elementary transformations. Even if this limits the possible choices for the transformation, it can be used as a first step, yielding a rough registration and, perhaps more importantly, yielding the correct feature pairs. The resulting feature pairs can then be used to find the parameters of more adequate transformations. While the method is reported to be robust against inconsistencies between features found in both images, it does present problems when too many features are available. In that case, the distribution can have several peaks or even become flat. In order to detect this type of problems the authors[12] defined two measurements of the quality of parameter estimate. A method of progressive feature filtering combined with a tree descent along peaks of progressively smaller importance is proposed to avoid these problems.

4.2.1. Feature Consensus for Finding the Parameters of the Similarity Transform

This transformation is characterised by four parameters (rotation angle β, translation t_x, t_y and scale s). Under this transformation a point x,y in the first image maps to the point X',Y' in the second image according to:

$$\begin{bmatrix} X' \\ Y' \end{bmatrix} = s \begin{bmatrix} cos\beta & -sin\beta \\ sin\beta & cos\beta \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} + \begin{bmatrix} t_x \\ t_y \end{bmatrix} \tag{3}$$

This transformation can be rewritten as a sequence of 4 stages:

$$\begin{bmatrix} X' \\ Y' \end{bmatrix} = T \begin{pmatrix} x \\ y \end{pmatrix} = T_{t_x} T_{t_y} T_s T_\beta \begin{pmatrix} x \\ y \end{pmatrix} \tag{4}$$

The first parameter to be estimated is the rotation angle. This can be observed from the slopes of line features in the image. If l and l' are corresponding line features with slopes ϕ and ϕ', we have:

$$\phi' = \phi + \beta$$

The rotation angle can thus be evaluated from the difference in line slopes. Please note that this does not allow to distinguish between rotation angles that differ by 180^o (i.e. β and $\beta + 180^o$). The correct choice has to be made either based on prior knowledge about the sensor platform's flight path or by backtracking from later stages of the transformation. Once the rotation angle has been found the rotation is applied to the objects from the map (or the first image) and then the scale factor is determined. This parameter is visible from the lengths of line segments or the distance between two point features. Only line segments that have a relative orientation that is approximately consistent with the already found rotation angle are considered in the voting. This limits the number of feature pairs to be considered while at the same time increasing the contrast between the peak of the votes and the random background. Limiting the number of votes to be considered at the various stages of the algorithm allows to increase the robustness of the method. Translation parameters are observed through point features. It is also possible to find the parameters of a semi-affine transform (i.e. different scale factors in x- and y-direction) using the feature consensus method.[12] In this case the rotation angle and scale factors are no longer visible separately. However, if the scale factors are close to 1, it is still possible to determine the rotation angles and scale factors separately which considerably facilitates the algortithm. This is the main reason why we rescale the SAR image to almost the same resolution as the map in the first step of the registration process. In fact, this even allows to use a similarity transform in the feature consensus function and to refine the resulting transform afterwards.

4.2.2. The Strategy used for the Feature Consensus Method

Our strategy for applying the feature consensus method involves three changes with the original method as presented in.[12] The first change is to distribute the vote of each pair of objects over neighbouring parameter values in the consensus function. The second is the use of 2D consensus functions for finding the scale factors (for the semi-affine transform) and the translation parameters. The third difference is the evaluation criterion that is used. Below, the three changes are described in detail.

Gaussian Filter of the Voting Functions

The characteristics of the primitives that are used for the feature consensus method are the mass-centres of the built-up areas and the forests. The extent of the built-up areas and thus also their mass-centre depends on the threshold applied to the results of the logistic regression in the detector for built-up areas. Furthermore the size of the built-up areas as well as that of the forests can change very quickly over time. It is therefore unlikely that the physical location of the mass-centres of these objects that is found on the SAR image corresponds precisely to the mass-centres of the same objects on the map. To take this imprecision into account, each primitive's vote in the consensus function is distributed over a range of transformation parameters. This is implemented by applying a Gaussian filter to the calculated parameter value in the consensus function.

2D Consensus Functions

In the original method[12] all parameters of the transform are determined sequentially, e.g. first the translation in the x- direction and then the translation in y-direction. However, it is better to couple some of the parameters. In case of the semi-affine transform for instance, if the scale factors are close to 1, the rotation angle can be determined independent of the scale factors.[12] Thus the rotation angle is determined and applied first and then line segments that are compatible with that rotation are used to first determine the scale factor along the x-direction and then the one along the y-direction. This gives two one-dimensional consensus functions for the scale factors. Fig. 3 (left) shows an example of such a set of one-dimensional consensus functions for the estimation of the scale factors of a semi-affine transform. An alternative is to couple the two scale factors in order to obtain a two-dimensional consensus function. The 2D consensus function, corresponding to the one-dimensional case is shown in fig. 3 (right). The 2D function offers several advantages. The 1D functions corresponds to the

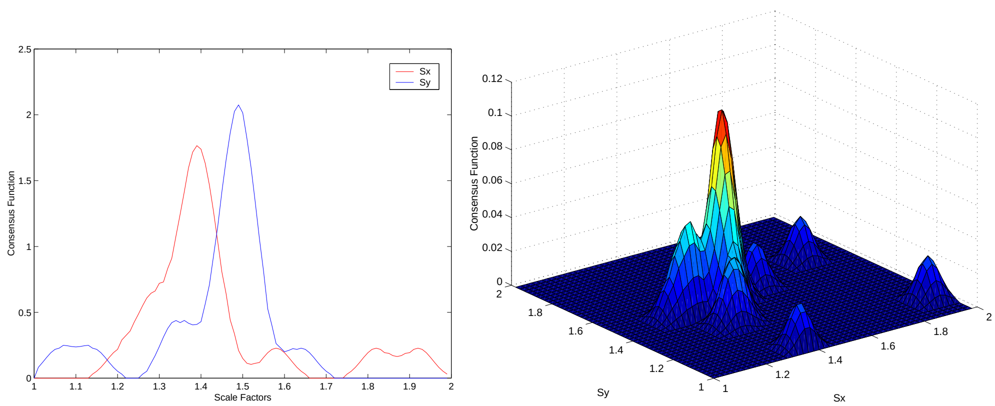

Figure 3. Consensus functions for the 2 scale factors: 1D functions (left), 2D functions (right)

projection of the 2D function in the $x = 0$ and $y = 0$ planes respectively. However, this projection will sum the peaks that are masked by each other along the projection direction. This means that two secondary peaks might merge to become the most prominent peak when the 1D consensus functions are used. The second reason to use 2D functions is that the two scale factors are coupled in a natural way, i.e. the best scale factor along the x-direction will only give a good transformation if it is coupled with the corresponding scale factor along the y-axis.

The Evaluation Criterion

We have tried to use the evaluation criteria that were presented in,[12] i.e. the peak distinctness and the parameter visibility, to guide the search trough the tree of possible transformation parameters. However their use does not lead to a robust method. Sometimes it is possible to find a clear peak for some parameter that is not the correct one. The only evaluation criterion that we found to be robust is the overlap between the objects detected on the SAR image and the objects from the map that are warped into the coordinates of the SAR image using the transformation that was found. This means that the complete depth of the search tree has to be explored before the evaluation can take place, i.e. all transformation parameters have to be determined before the overlap can be evaluated. In order to determine the overlap, the map objects are warped onto the SAR image. Let G_M be the set of pixels corresponding to the interior of a warped object from the map and falling inside the SAR image. G_S is the set of pixels in the interior of the same type of object detected in the SAR image. The overlap quality factor Q_o is then defined as:

$$Q_o = \frac{Card(G_M \bigcap G_S)}{Card(G_M \bigcup G_S)}, \tag{5}$$

with $Card$ representing the number of points in a set. If various types of objects are used in the feature consensus method, for each of them the quality factor is determined and the average quality factor over all objects is used as an evaluation criterion. The feature consensus method selects the set of transformation parameters which gives the highest value of Q_o.

Finding the Maxima

As the number of inconsistencies or differences between the set of objects found on the map and in the image increases, several peaks may occur in the distribution of the different parameters. It is not necessarily the highest peak that is correct. Therefore it is not sufficient to find the highest value of the consensus function for a given parameter, we need to find all local maxima. Starting from the local maximum with highest value, the transformation parameters are determined and the registration is evaluated. In order to find the local maxima of the parameter's histogram we used an approach based on morphological operators. The method is represented in figure 4 for a 1-dimensional case. In the figure, "H" is the morphological structuring element.

4.2.3. Application of the Feature Consensus Method to the SAR Image

As already mentioned, the feature consensus method is used to find a first rough registration between the SAR image and a topographic map. The matching primitives that are used here are the location of built-up areas and/or forests. In fig. 5 the objects detected on the SAR image are shown at the left and the part of the map corresponding to the SAR image on the right.

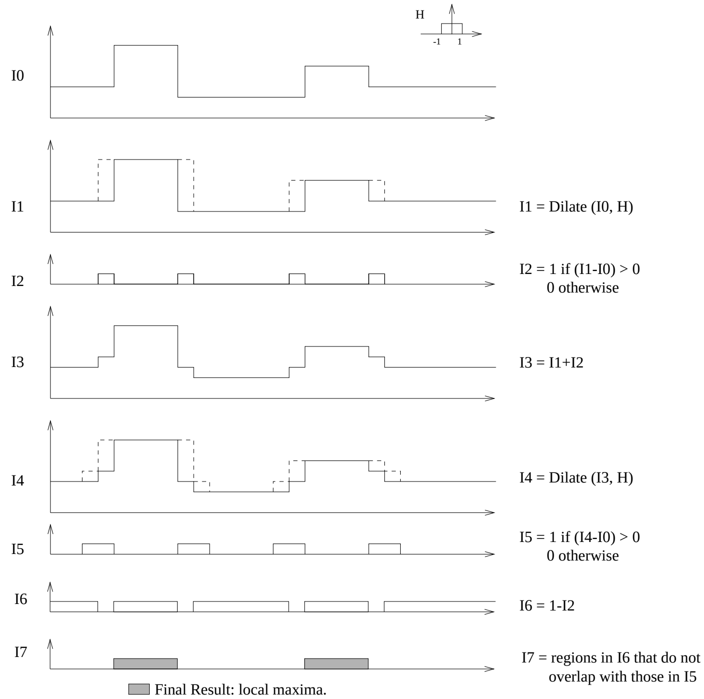

Figure 4. Detection of local maxima in the consensus function

Figure 5. Objects Detected in SAR Image and Region on the Map corresponding to the SAR Image

It is possible to constrain the search in the feature consensus method by limiting the range of possible transformation parameters. This can be done on the basis of knowledge of the SAR image acquisition parameters (i.e. heading, pixel spacing and GPS coordinates). In order to test the robustness of the method we have set the ranges for the different parameters very wide (rotation angle: $\beta \in [0..360^o]$, scale factor $s \in [1..3]$, translations $t_x, t_y \in [-3000..3000]$). Because the pixel spacing of the SAR image is always approximately known, it is reasonable to set the range for the scale factor somewhat narrower than for the other parameters. The feature consensus method was applied to search for a similarity transform. The resulting matched objects are then used in a first refinement step to find the parameters of an affine transform using a least square method. The left image in fig. 6 shows the results of the feature consensus method and the subsequent refinement, based only on the position of built-up areas. The image shows the objects detected in the SAR image as well as the objects from the map that were warped to the SAR image with the found transformation.

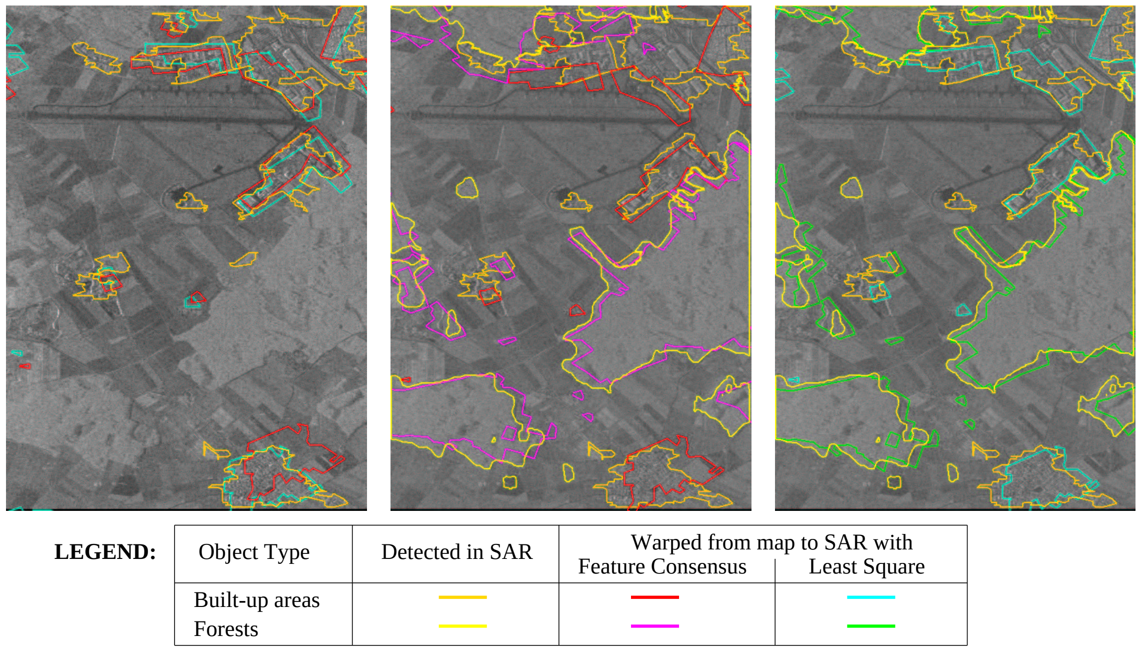

Figure 6. Result of feature consensus method using only built-up areas (left) and using both Built-Up and forests (middle: using feature consensus method and right: using least square method).

The two other images in figure 6 are the results of using both the built-up areas and the forests. In the middle the similarity transform found directly by the feature consensus method is applied. On the right the results of applying the affine transform are shown. Contrary to what could be expected, the result of the feature consensus method using only the built-up areas is slightly better than when the positions of the forests are used too. This is due to the fact that several of the forests cross the border of the image which results in a localisation of the centres of the forests that is different in the image and the map. These differences widen the peaks in the consensus function which makes the parameter estimation less precise. However, in both cases the objects are matched correctly and using the least square method to find the parameters of an affine transform does result in a comparable registration. In fact, the result after the least square method is slightly better when both built-up areas and forests are used. These results of visual comparison are confirmed by the evaluation method described earlier (see table 1).

Input	Method	Transform	D_{mean}	D_{RMS}	D_{max}
Built-Up Areas	Feature Consensus	Similarity	31	35	63
	Least Square	Affine	26	33	71
Built-Up Areas and Forests	Feature Consensus	Similarity	49	52	75
	Least Square	Affine	17	18	35

Table 1. Evaluation of the Registration Results

4.3. Refinement of Registration using Contours of Forests

The contours of forests from are approximated by polygons. The transform found by the feature consensus method is used to warp the contours from the map into the SAR images. Then nearby vertices of the "map" and "SAR" polygons are used to refine the affine transform using a least square method. This process is iterated until the parameters of the affine transform remain unchanged.

4.4. Refinement of Registration using Communication Lines

The forests and built-up areas can only provide a first registration because their 3D structure results in an inaccurate localisation. The obtained registration should be further refined using features with a low 3D structure. Good candidates for refining the registration between maps and images are roads, railroads and waterways. This communication network, extracted from the map, is warped onto the SAR image using the affine transformation that was found so far. The result (fig. 8 (left)) shows that detecting the communication network on the SAR image is indeed likely to provide a possibility to improve the registration further.

The detector for communication lines gives a set of line segments that potentially correspond to roads or waterways (it detects "dark bars"). The result of this detector can thus be used for refining the registration. The proposed refinement procedure relies on the following hypotheses:

- the current registration is already relatively accurate. This means that candidate road segments only need to be searched in a restricted neighbourhood of the roads warped from the map into the SAR image.
- as we have detected only segments of roads in the SAR image, a matching between the roads from the map to these segments will only yield the displacement component perpendicular to the road.
- we search for a global affine transformation

The first hypothesis allows to match most of the line segments found on the SAR image to one or more roads found on the map. The matching can be restricted to a small neighbourhood and to lines that are almost parallel. The second hypothesis restricts the methods that can be used to find the transformation. Please note that the method presented below was applied directly to the results of the feature consensus method refined with a least square approximation of the affine transform but using only the positions of built-up areas. The idea is that forests could be considered as features that are not stable over time, i.e. forests can disappear or their surface area and contour can change very quickly. The affine transformation can be written as:

$$\begin{bmatrix} X' \\ Y' \end{bmatrix} = \begin{bmatrix} T_x \\ T_y \end{bmatrix} + \begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} X \\ Y \end{bmatrix} \tag{6}$$

Applying the affine transform with a given set of transformation parameters $\Theta = T_x, T_y, a, b, c, d$ will transform a point (X, Y) into a point (X', Y') and the displacement of the point P : $\overrightarrow{dP} = dX\overrightarrow{1_x} + dY\overrightarrow{1_y}$ is given by:

$$dX = T_x + aX + bY - X \tag{7}$$

$$dY = T_y + cX + dY - Y \tag{8}$$

If a matching between the line segment from the SAR image and a road from the map can be found, we know the component, perpendicular to the road on the map, of the displacement that the points of the SAR line segment should undergo, when the optimal transform is applied. The aim of the optimisation method is thus to find a set of transformation parameters that yields a displacement that minimizes this perpendicular component, i.e.

$$MIN_{\Theta} \left(\sum_i (\overrightarrow{dP_i}\overrightarrow{1_{D_\perp}} - D_\perp(P_i) \right)^2 \tag{9}$$

with i an index over the points of the SAR image for which a match with the map was found. $D_\perp(P)$ is the value of the perpendicular distance and $\overrightarrow{1_{D_\perp}} = (1_{Dx}, 1_{Dy})$ a unit vector pointing along the perpendicular direction from the point of the SAR segment to the road from the map (see fig 7).

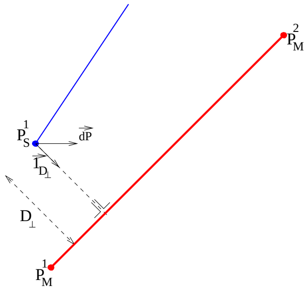

Figure 7. Principle of optimisation by least square method

The minimisation problem can be written as a system of equations:

$$[1_{Dx}]\ T_x + [X\ 1_{Dx}]\ a + [Y\ 1_{Dx}]\ b + [1_{Dy}]\ T_y + [X\ 1_{Dy}]\ c + [Y\ 1_{Dy}]\ d - [X\ 1_{Dx} + Y\ 1_{Dy} + D_\perp(P)] = 0 \quad (10)$$

The unknowns $\Theta = T_x, T_y, a, b, c, d$ are found by a least square method. Although the least square method should yield the optimal set of parameters in one iteration, applying the method iteratively improves the results. This is due to the fact that in the first run a number of matches can be false. These pull the solution away from the correct one. It is not possible to identify these incorrect matches a priori but, as with the feature consensus method, the incorrect matches are more or less randomly distributed while the correct matches all induce a preferred displacement corresponding to the correct transformation. The least square method will thus, in the first step, approximate the solution. After applying the transformation that is found, some incorrect matches will become invalid and in the next iteration the resulting registration accuracy will thus be improved. Table 2 shows the distance to the ground control points at successive iterations. The line corresponding to 0 iterations is the result of the feature consensus method using only built-up areas. The result of warping the communication lines on the SAR image are shown in fig. 8 (right). In the figure the results of another optimisation method (the simplex method[13]) are also shown. Although, in this example results obtained by both methods are similar, tests on other images have shown the least square method to be more robust. The simplex method is therefore not explained here.

Iteration	D_{Mean}	D_{RMS}	D_{Max}	Iteration	D_{Mean}	D_{RMS}	D_{Max}
0	26.0	33.5	71.0	4	18.1	20.5	42.4
1	23.0	28.6	62.7	5	13.8	15.7	33.1
2	22.8	27.3	59.7	6	8.5	9.5	20.2
3	21.2	24.2	48.5	7	6.8	7.5	15.1
4	18.1	20.5	42.4	8	5.6	6.5	13.9

Table 2. Distance between GCP's as a function of the number of iterations for the least-square method

5. REGISTRATION BETWEEN DIFFERENT IMAGES

Once each image is registered with a map, the registration between different images becomes easy: The registration between the images and the map is used to obtain a first approximation of the affine transform between the different images. That affine transformation is refined using the position edges detected in the two images. The edge detectors were applied to the original slant range images and therefore the resulting objects need to be first converted to "ground range" and rescaled in the same way as the SAR images. In fig. 9 a part of a second SAR image is shown. In blue the edges detected in that image are represented. The yellow objects correspond to edges detected in image 1 and warped into the second image using the affine transform that was found so far. The two sets of edges represented in the figure are used to refine the registration. In fig. 9 (right) the results of the refinement are shown. For the refinement the same least square method is used as for the refinement

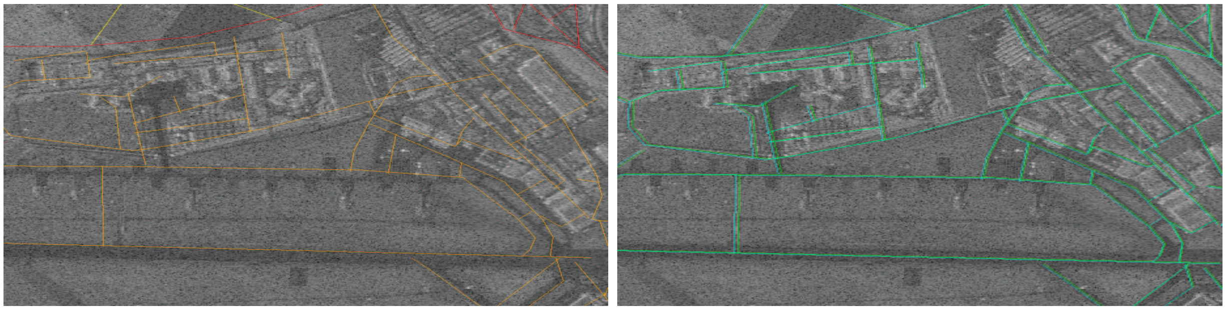

Figure 8. Communication lines from the map warped to the SAR image before refinement (left) and after the refinement using the Simplex Method (Green) and the LSQ method (Cyan)

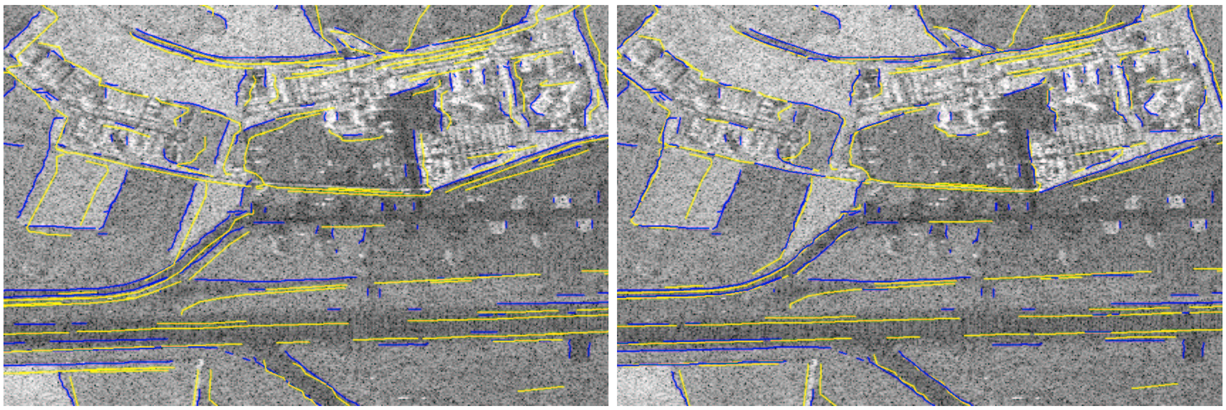

Figure 9. Detected and warped edges before and after the final refinement

based on communication lines. In order to get a quantitative idea of the accuracy of the obtained registration between the two images, the evaluation method described before can be used. However, the evaluation method uses manually determined ground control points. For a rough registration, the error made by manually selecting ground control points such as road crossings can be neglected with respect to the registration error. Now, after the final refinement step, this is no longer the case and we need to find control points that are the same in both images to within an accuracy of one pixel. The only possible candidates are the corner reflectors present on the airfield. The control points are the pixels of highest intensity of each of the corner reflectors. The results of the evaluation are shown in table 3. The first line in the table is the result of the registration only using the combination of the registration of both images with the map. The second line presents the results after refinement using edges.

Method	D_{mean}	D_{RMS}	D_{max}
Before refinement	2.428	2.523	3.299
After refinement	1.018	1.106	1.726

Table 3. Evaluation of the Final Registration Results

6. CONCLUSIONS AND FURTHER WORK

In this paper a method for registering multi-sensor images for remote sensing is presented. Features that could be used for registering different types of images are edges because they usually correspond to boundaries of physical objects on the scene. However, in many cases, they are not discriminative enough to be directly used for registration, i.e. although edges of agricultural terrains could allow an accurate registration if correctly matched, it can be difficult to match the correct edges because neighbouring fields can be very similar. Edges can only be used in a refinement step. In this article an hierarchical registration method is proposed that uses topographic maps as a help in the registration. The basic idea is that for any type of image it should be possible to identify objects that are present on the map and easily detectable on the image. The resulting registration between each image and the map can then be used to derive a first registration between the different images. This registration is then refined based on features that are visible on all images. Edges (or uniform regions) are an obvious choice as they usually correspond to physical objects in the scene. The method was applied for the registration of polarimetric SAR images. For such images a first registration between image and map can be obtained based on the location of built-up areas and/or forests. This registration can then be refined using communication lines. Although the proposed method was applied for registering polarimetric SAR images, the method can be used for any type of image; only the object/feature detectors need to be adapted. Until now we have only used a global affine transform to register the (ground-range) SAR images. This was possible because the imaged terrain is very flat. We plan to refine the registration further using local transformations and available digital elevation models.

ACKNOWLEDGMENTS

The presented work is the result of a collaboration between the Signal & Image Centre of the Belgian Royal Military Academy and the Institut für Hochfrequenztechnik und Radarsysteme of the Deutshes Zentrum für Luft- and Raumfahrt (DLR), who also provided the images. The images used in this project are full-polarimetric L-band SAR images acquired by DLR's E-SAR system.

REFERENCES

1. P. Dare, R. Ruskoné, and I. Dowman, "Algorithm development for the automatic registration of satellite images," in *Proc. Image Registration Workshop*, NASA, (Goddard Space Flight Center), November 1997.
2. P. Dare, *New Techniques for the Automatic Registration of Microwave and Optical Remotely Sensed Images.* PhD thesis, Dept. of Geomatic Engineering, University College London, London, April 1999.
3. D. Borghys, C. Perneel, and M. Acheroy, "Detection of built-up areas in polarimetric sar images," in *1st Int. Workshop on Pat. Recog. in Remote Sensing; Andorra*, sept 2000.
4. D. Borghys, C. Perneel, and M. Acheroy, "Automatic detection of built-up areas in high-resolution polarimetric sar images," *Pattern Recognition Letters: Special Issue on Pattern Recognition for Remote Sensing* , 2001.
5. S. Cloude and E. Pottier, "A review of target decomposition theorems in radar polarimetry," *IEEE-GRS* **34**, pp. 498–518, March 1996.
6. M. Hellmann, *Classification of Fully Polarimetric SAR-Data for Cartographic Applications.* PhD thesis, Fakultät Elektrotechnik der Technischen Universität Dresden, Dresden, Juni 2000.
7. D. Borghys, C. Perneel, and M. Acheroy, "A multi-variate contour detector for high-resolution polarimetric sar images," in *International Conference on Pattern Recognition, Barcelona*, vol. 3, pp. 650–655, sept 2000.
8. D. Borghys, C. Perneel, and M. Acheroy, "Contour detection in high-resolution polarimetric sar images," in *SPIE Conference on SAR Image Analysis,Modelling and Techniques; Barcelona*, sept 2000.
9. F. Tupin, H. Maitre, J. Mangin, J. Nicolas, and E. Pechersky, "Detection of linear features in sar images: Application to road network extraction," *IEEE-GRS* **36**, pp. 434–453, March 1998.
10. C. Shekhar, V. Govindu, and R. Chellappa, "Multisensor image registration by feature consensus," tech. rep., University of Maryland, Computer Vision Laboratory, Center for Automation Research, August 1996.
11. V. Govindu, C. Shekhar, and R. Chellappa, "Using geometric properties for correspondence-less image alignement," *IEEE* , pp. 37–41, 1998.
12. C. Shekhar, V. Govindu, and R. Chellappa, "Multisensor image registration by feature consensus," *Pattern Recognition* **32**, pp. 39–52, 1999.
13. W. Press, S. Teukolsky, W. Vettering, and B. Flannery, *Numerical Recipes in C*, Cambridge University Press, Cambridge, 1992.