Reconstruction of planar surfaces behind occlusions in range images

Fabio Dell'Acqua, Robert Fisher
Division of Informatics, University of Edinburgh
Email: {fabiod,rbf}@dai.ed.ac.uk

Abstract

Analysis and reconstruction of range images usually focuses on complex objects completely contained in the field of view; little attention has been devoted so far to the reconstruction of simply-shaped wide areas like parts of a wall hidden behind furniture pieces in an indoor range image. The work presented in this paper is aimed at such reconstruction. First of all the range image is partitioned based on depth discontinuities and fold edges. Next the planes best fitting each of the regions constituting the partition of the image are determined. A third step locates potentially contiguous surfaces, while a final step reconstructs the hidden regions. This paper presents results for reconstruction of the shape of planar surfaces behind arbitrary occluding surfaces. The system proved to be effective and the reconstructed surfaces appear to be reasonable. Some examples of results are presented from the Bornholm church range images.

Keywords: image processing, occlusion, range data analysis, range image partition, range data reconstruction.

I. Introduction

Range images are used in a wide range of applications. So far they have been extensively used in object recognition [10][15][8][9], reverse engineering [6], and other applications, nearly all focused on small and rather complex objects. While extending the use of range images to a whole environment rather than well-delimited objects (an important example of these applications is the CAMERA EU project [5]) new issues arose. Occlusion is a major cause of information loss: with even moderately complicated objects it is virtually impossible or impractical to obtain complete range scans. Still, an exhaustive description of the observed object or environment is needed for some applications, like construction of a 3-D model [5]. An alternative way of filling in the gaps, at least partially, without performing extra scans, is to infer the layout of objects in the occluded area by exploiting information from the surroundings. This procedure is termed reconstruction. Reconstruction of simple-shaped wide regions occluded by objects located closer to the sensor is an important point for this procedure. In this paper we propose a method to reconstruct plane surfaces behind furniture pieces or other closer objects. This is a new problem, on which little previous work [7] has been done. In this paper, section II presents the data we used for our experiments, section III presents the theory, section IV presents the experiments and discussion, while section V draws some conclusions.

II. The data

Range information can be obtained by a variety of methods [1]. The data we used was sensed by a K2T 3-D laser scanner on 360 degrees azimuth, 63 degrees elevation, both with a 0.045 degrees step, thus resulting in 8000×1400 pixel images. For each pixel both range and IR intensity values are registered. Range precision is around one millimeter; range values are converted into x,y,z values in a cartesian system centered on the sensor.

The huge dimensions of the raw range data files (hundreds of MB) pushed us to work on small subimages, i.e. rectangular subsets of pixels.

In figure 1 on the left a subimage extracted from a range image is shown. Note that a planar surface is clearly visible lying behind the chair, and a part of the surface is hidden by the chair itself. The aim of our work is to reconstruct such missing data.

III. Theory of Occluded Surface Recovery

The key to reconstruction is to identify contiguous surface regions potentially connected behind closer occluding surfaces (identified by depth discontinuities and closer distances). Hypothetical surfaces can then be created to connect or fill in the contiguous surfaces behind the occluding surfaces.

A. Planar image region segmentation

Preprocessing with a median filter [13] reduced the 4.62% outlying pixels. The next step is to partition the image into different surface regions that may belong to the same surface, although separated by an occluding object. Separation between adjacent regions can be marked [11] either by a depth discontinuity (e.g. a boundary of an occluding object) or by a fold edge, which is by definition an area where the local surface normal changes suddenly (e.g. where two walls meet). Other cases have not been considered [19]. To reduce the complexity of the general problem of segmenting a range image [2], [14], [18], this was split into two stages. The first part finds depth discontinuities; then, a further partition based on fold edges is performed. These two procedures are described in the following.

A.1 Depth discontinuity detection

Depth discontinuities in closely sampled range images are unexpectedly smooth, probably due to the aperture of the laser beam. In our data, the shortest across-discontinuity path usually averages between three and eight pixels. A simple threshold on gradient is thus unsuitable for partitioning the image, and we chose to use a more refined Canny edge extractor. First we build three 2D images out of the x, y, z co-ordinates for each pixel. Then we apply the Canny edge extractor to each of those images and combine the results by means of a three-input "OR" operation applied to each pixel. As a discontinuity cannot be parallel to all three of the axes, wherever a gap in the co-ordinates is present, it will be detected in at least one of the images. The final "OR" operation fuses all the partial results together. Call σ the standard deviation of noise on point coordinates, which turned out to be around 7mm. Then, lower and higher Canny thresholds were set to about $\frac{1}{2}\sigma$ (3 mm) and 3σ (20 mm) respectively. This is to make false activation extremely unlikely (noise$> 3\sigma$). The next step is a double dilation on this edge image, aimed at two purposes. The first is to include in the "edge region" (pixels set to 1, ignored in the computation of region features) all the pixels around the edge, likely to be corrupted by the "edge smoothing effect" described above. The second purpose is to help bridge possible, though infrequent, small gaps in the image partition. At the end of this stage the binary image partitions the original image into regions. Each region is extracted and labelled. An example result is shown in figure 1 on the right.

Fig. 1. A range image ($left$) and corresponding extracted regions after labelling($right$).

A.2 Fold region selection

Some of the regions extracted using the procedure described in Section III-A.1 may contain fold edges. In this section we discuss how to detect them. The best fitting plane is computed for each region (the method used is described in Section III-B), and the mean square distance σ_d of the points from the best fitting plane is considered. The more the shape of the region departs from a plane, the higher the resulting σ_d. A threshold of 15mm has experimentally been found suitable to discriminate regions possibly containing fold edges (higher σ_d) from the nearly-flat ones (lower σ_d). Only the former undergo the fold edge extraction procedure described in the following.

A.3 Fold edge detection

Fold edges in range images are difficult to detect due to the high level of noise which prevents reliable determination of a local surface normal vector. A number of methods have been proposed for estimating surface normals, but none outperforms the others, and the choice of the method is not trivial [12]. We considered the methods presented in the overview contained in [12] and combined the features of the various methods in order to address the problems we encountered, in particular the high noise level. To get an acceptable estimate of the surface normal it is necessary to renounce high resolution and operate on wide-area averages. The x,y,z image with outliers removed is first low-pass filtered to remove more noise and then partitioned into a grid of 5×5 pixel patches. On every patch, the plane best fitting its points is determined. Then, two patches are considered matching when the cosine of the angle between the two normals is higher than 0.9 (26 degrees) and the difference between their distance from the origin is lower than 0.2m (20σ, σ=standard deviation of noise values). Both values were determined by experiment. For every patch, the number of matching patches are counted. After that, the patch with the highest number of compatible patches makes the first partition; the procedure is then repeated using only the remaining patches, again picking the largest set of matching patches and so on; the procedure stops when no new set bigger than 20 patches can be formed; smaller sets generally produce unreliable best fitting planes. Any remaining patches are marked as non-usable. Isolated patches and smaller non-connected components are then suppressed to regularise results. Eventually, results are combined

with the previous depth-discontinuity partition to give final image partition. To illustrate the effectiveness of this technique, consider the sample subimage in figure 2. It contains a large fold edge on the wall corner at the right, as well as several smaller ones on the steps. Note the roughness of the wall surface which, together with the high level of high-frequency noise, makes extraction of local normal vectors unreliable.

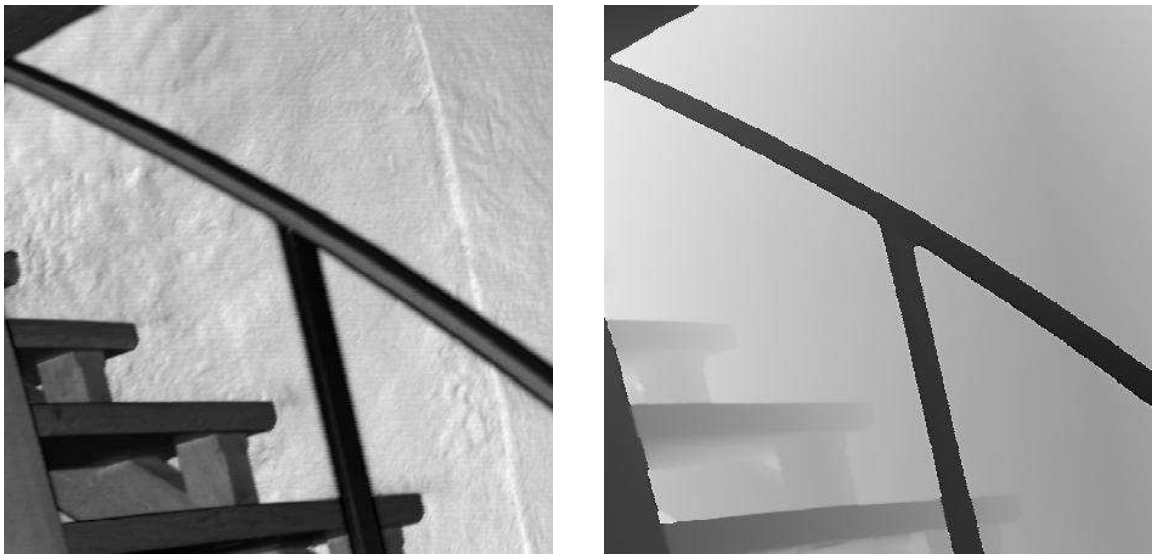

Fig. 2. An intensity image (*left*) and its corresponding range image (*right*).

Figure 3 shows the image partition. The image on the left displays the partition based on simple depth discontinuity detection. The image on the right shows image partition after fold edge detection with the two sides of the wall separated, as well as some of the many planar surfaces on the steps.

B. Region description

Once the image has been partitioned, for every region the best fitting plane is determined by means of a least squares approach. The least squares method provides a three dimensional unit normal vector $\vec{n}$ (directed towards the observer) and a scalar representing the distance of the plane from the axes origin. These two parameters are recorded and associated with the analyzed plane.

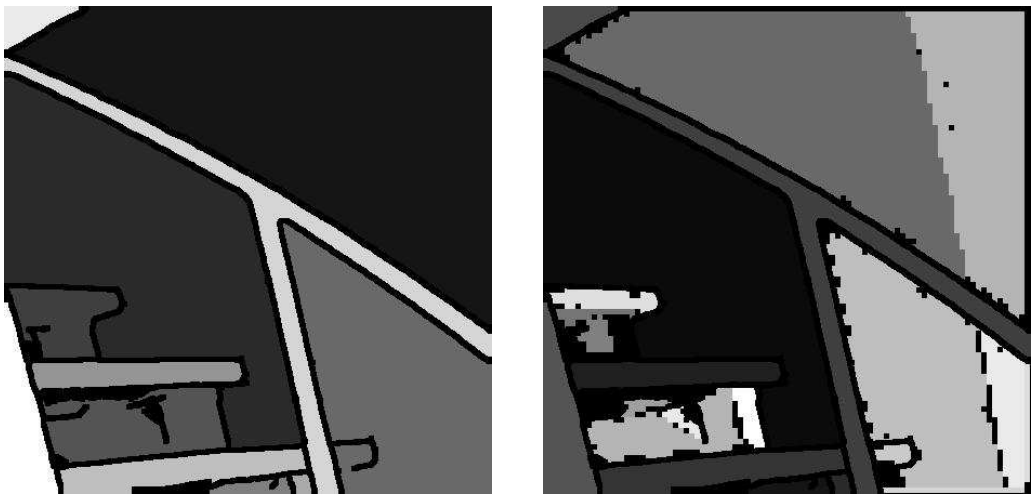

Fig. 3. The image in figure 2 after some processing. Left: partitioned using only depth discontinuities. Right: additionally partitioned using also fold edges.

C. Region reconnection

The search for possibly contiguous, occluded regions is next. Two regions are deemed possibly contiguous if the scalar product between their normals is higher than 0.92 (corresponding to normals lying less than 23 degrees apart) and the difference between their displacements with respect to the origin is less than 0.25 m (around 35 times σ of noise). Within an image all the possible region pairs are compared and organized into groups. Groups are such that each region within a group matches with each other region within the same group, according to the definition given above. The constraints on normal orientation and plane distance from the origin are necessary for two regions to be reconnected. However, they are not sufficient, as two separated areas lying approximately on the same plane can be separated by an object that is further than themselves; in this case, no occlusion of the plane is taking place in the area in between and reconstruction makes no sense. An example may be a niche separating two parts of the same wall.

Thus, we need to test the connecting regions between, to decide whether occlusion has occurred. We reconstruct potentially occluded areas by a pairwise merging of visible regions across occluded regions, and we test hypothesised occlusion by investigating the area between the boundaries of the two regions. The process starts with the two lists of the boundary pixels from each region in the pair under test. Then each possible pair of boundary pixels, one from each list, is considered. The Bresenham algorithm [4] is used to

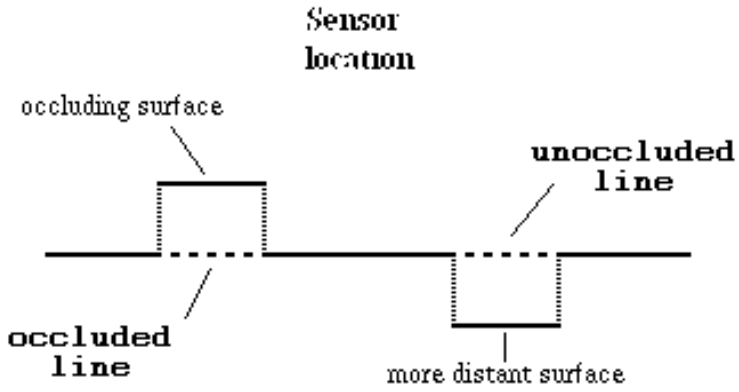

Fig. 4. Occluded and unoccluded surface patches. The solid line represents the observed surface, while the dashed lines represent the test paths.

generate pixels on the path from the first region boundary to the second region boundary. By interpolating the depths of the endpoint pixels, a depth is associated with each pixel on the path. Each interpolated depth is compared with the actual sensed depth. If the interpolated depth is closer than the observed depth it is evidence against the possibility of occlusion, and the event is recorded. To take into account the high noise level we had to allow a margin of 4σ on the observed point distance from the plane best fitting the surface. More precisely, the pixel is considered occluded if:

$$d_n < d_1 + \frac{(d_2 - d_1)}{N} \cdot n + 4\sigma \tag{1}$$

where d_1 and d_2 are the depths of the line endpoints, N is the number of pixels separating the endpoints, σ is the standard deviation of actual point distances from the best fitting plane. d_n is the actual range data for the pixel currently being considered, which is the n-th of N. A line is classified as unoccluded if none of its points have evidence against occlusion. This means that the pixel set covered by the line has a high probability of belonging to an object generating an occlusion. This may be better understood by looking at the scheme in figure 4.

The dashed line on left of the picture represents an occluded surface: from the observer the sensed surface (solid line on top) is closer than the line interpolating the surrounding regions. The dashed line on the right is non-occluded, as interpolation between surrounding

regions is closer than the observed surface (solid line on the bottom).

As the number of boundary pixels is generally on the order of hundreds or thousands, the number of possible endpoint pairs can be huge and a selection step is needed to keep processing time within reasonable limits. A first elimination criterion relies on the fact that it is useless to test lines crossing the considered regions themselves. More ways to reduce the computational effort are to consider only the shorter lines from one region to another, and to limit the number of lines starting from each point.

After the connecting lines have been tested, the number of occluded lines (lines reporting points closer than the local threshold) is tested for consistency with the hypothesis of occlusion. In theory a single occluded line might connect two surfaces, but this would require a very thin connection and is thus unlikely. Thus, two regions are considered to be separated by an occluding object if the number of occluded lines exceeds 10.

D. Reconstruction

Reconstruction takes place between regions belonging to the same group and satisfying the depth constraint. Groups may contain more than two regions, but for sake of simplicity, explanation is first given for a single pair of regions and then extended to the general case. Given two matching regions, reconstruction only takes place on pixels belonging to lines satisfying the occlusion path constraint described in Section III-C. This is to ensure that reconstruction will not:

1. involve areas that do not satisfy the maximum depth constraint;
2. extend to areas too far from the nearest known data.

As any reconstruction requires hypothesising non-observable data, these assumptions lead to a conservative reconstruction having a high likelihood of being correct. While testing lines, a map is generated out of all the pixels belonging to accepted lines; these pixels undergo reconstruction. An example of such map, still using the example image in figure 1, is shown in figure 5, left. The white region contains occluded lines.

Where all the tests have given a positive result, reconstruction is performed. Every pixel in the reconstructed area is assigned a new depth by interpolating the depths at the end points of the shortest line passing through the pixel and joining two boundary pixels in matching regions. As a result, the occluding object seems to have been removed from the

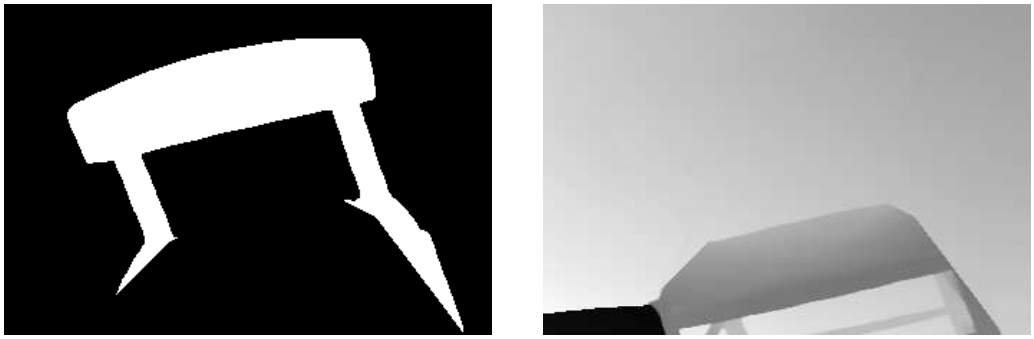

Fig. 5. Left: map of occluded lines, in white on a black background. Right: range image after occluding points have been replaced by estimated occluded points(compare with figure 1, left)

scene, as visible in the corrected range image in figure 5, right.

However, further analysis using a cosine shaded image like in figure 6 on the left can give a better idea of the situation. The reconstructed area is indeed still distinguishable from the true wall, even though it is not so evident. This is due to wall texture, which is missing in the reconstructed area; simple interpolation cannot reproduce it satisfactorily, especially for large areas.

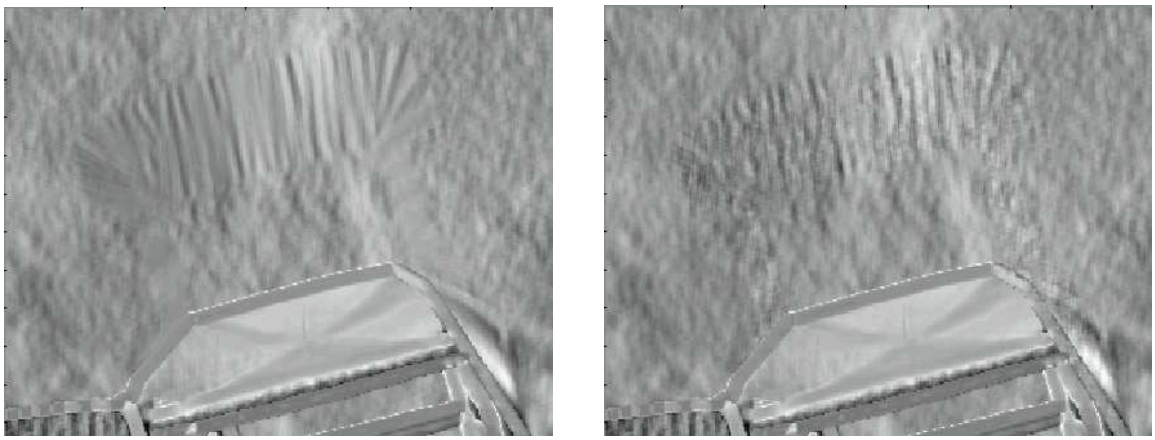

Fig. 6. Left: the cosine shade image for reconstructed version of the subimage in figure 1. Right: same as above, with noise added.

We make the reconstructed area more homogeneous with the original regions by estimating the noise level in them and adding a similar amount of noise to the reconstructed area. The improvement is visible in figure 6 on the right.

Still, if only a rough reconstruction of the area behind occlusion is required, the algorithm does a good job of filling gaps. This can be noted in figure 7, where 3-D represen-

tations of the scene are displayed, before and after reconstruction.

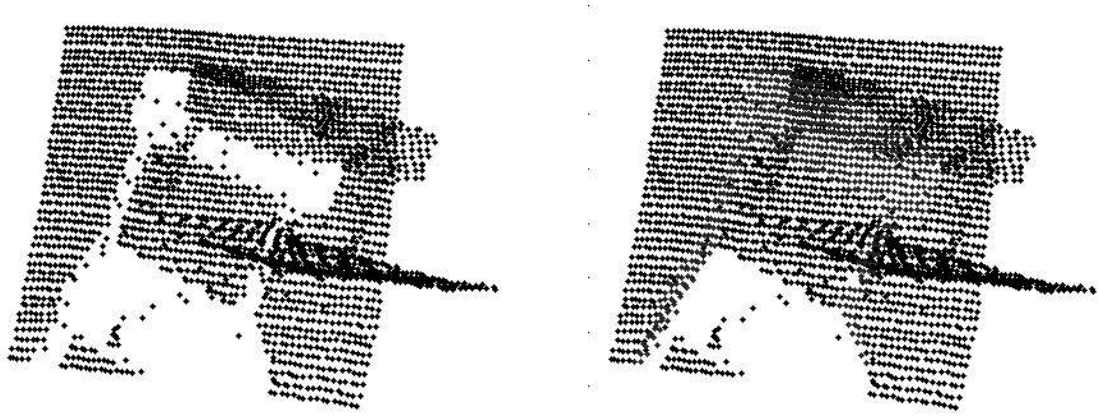

Fig. 7. Original x,y,z image (left) and the same image with the reconstructed pixels added (right).

IV. Some results

The algorithms described in Section III were implemented in Matlab and run without compilation on a Sun Ultra 10 machine. The process takes about half an hour per image on average, though duration is dependent on image complexity. The largest contributions to total processing time comes from fold edge extraction and planar approximation. A substantial improvement is expected from recoding the process and compiling the scripts, possibly resulting in cutting the time per image to less than one minute.

We used the Bornholm Church range data for more experiments. From the 9 original indoor range images a set of 19 small subimages were extracted. Sizes of subimages ranged from 231×151 to 751×501 pixels. The depth discontinuity method (Section III-A.1) first partitioned the subimages into 90 regions. Then, the fold edge method (Section III-A.3) further partitioned the 90 regions into a total of 160 regions, an average of 9.26 per subimage. The least partitioned subimage contained only two regions, while the most partitioned one had 22 regions.

In 2 subimages the fold edge method increased the number of regions from 1 to 10. Such subimages depict niches with no depth discontinuities where segmentation can only rely on fold edges. In 3 subimages the fold edge method caused no further partition as they actually contained no fold edges.

Upon completion of subimage partition, a human observer marked matching regions

("ground truth"), exploiting the clear visual information contained in corresponding intensity images. Within the set of 946 possible region pair combinations, only 46 pairs are actually matching.

	Matching pairs	Non-matching pairs
Declared potentially matching	41	120
Declared non-matching	5	780

Confusion matrix for region pair matching, before occlusion test

Note the imbalance between false matching (12.7%) and missed matching (0.53%) rates. We chose to keep coplanarity requirements loose to minimise missed matching, relying instead on the occlusion test (Section III-C) to reduce the final rate of false matching.

The 5 missed matching cases were inspected. The regions involved are small (<500 pixels) and have an elongated shape. Border effects thus corrupt results to a wider extent than with regular data; moreover, the slope of best fitting plane is more easily biased by noise.

Based on reciprocal matching, the 97 regions involved in the 161 matching pairs were collected into 31 groups containing an average of 3.12 regions each. Of the 161 pairs, only 78 passed the occlusion test and were thus deemed suitable for reconstruction. This table shows the final classifications:

	Matching pairs	Non-matching pairs
Declared matching	41	37
Declared non-matching	5	863

Confusion matrix for region pair matching, final results

Inspection of the images showed that the 41 correctly declared matching regions (52.6%) were actually suitable for reconstruction, while the remaining 37 were not. This latter figure seems quite high, but please note:

1. Included are the cases of "useless" reconstruction; e.g., where partition is due to noise. Correct plus useless reconstruction score 46 (59.5%) together;
2. In any case, reconstruction only takes place where occlusion is possible, and in most cases this is sufficient to reduce to a minimum the effect of a wrong choice. The number of invalid pixels generally ranges from tens to a few hundreds.

A nullified incorrect matching (see point 2 above) is visible in the following example. In figure 8 on the left, the intensity image for one of the subimages is shown.

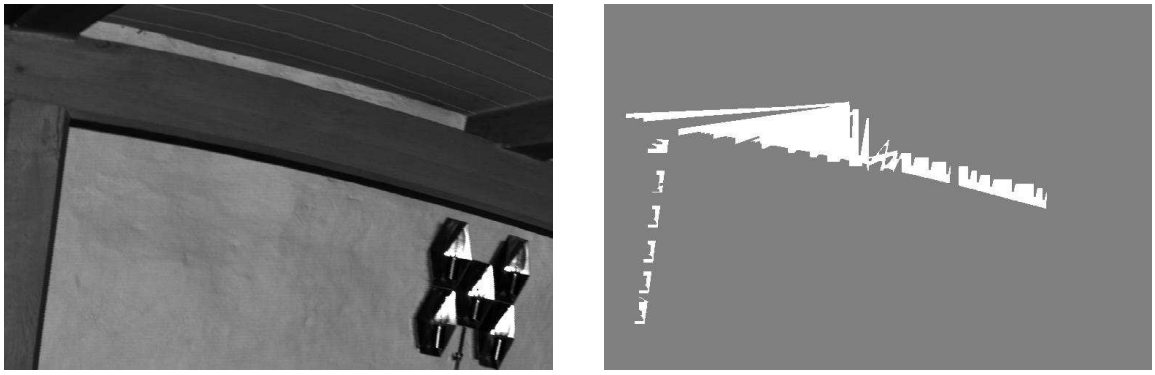

Fig. 8. An example subimage ($left$) and the corresponding set of reconstructed points ($right$), in white on black background.

Fold edge partition created two regions on the faces nearest to the observer of the two dark beams parallel to the wall, due to noise in between. One region includes the vertical beam, while the other region includes the rightmost part of the horizontal beam. Coplanarity of the faces made this pair pass the first test. The second test (minimum number of occluded connecting lines) also passed because of a group of lines running from the top of the vertical beam to the nearest part of the horizontal beam, on the flat surface at the crossing of the two beams. The two regions were thus deemed suitable for reconnection. Additional tests could remove some of these cases.

However, thanks to local occlusion test, reconstruction took place to a negligible extent, mostly flattening small bumps on the beams. The wall was untouched, as seen in the reconstruction map in figure 8 on the right, where the white area is the set of reconstructed points. Indeed, replacing real data points with points from the best fitting plane meant little change or no change at all. Original and reconstructed 3D point sets are nearly undistinguishable.

Still, some remarks are to be made on the cause of false matching. A brief investigation revealed that the main reasons for such errors were:
1. Random matching of small regions, whose parameters are heavily biased by the noise;
2. Forced partition of concave curved regions into planar patches. Partition of a curved region into planar patches creates a range of slopes and depths available for matching, while concavity provides surrounding regions satisfying the occlusion constraint on test paths. A typical example is the back of a chair against the wall, which accounts for the majority of the false matching cases reported. The result of such a false matching is a surface that connects a portion of the chair back to the wall.

Where applicable, results of reconstruction are hard to judge quantitatively, because no truth is available for comparison. From a qualitative point of view, results seem reasonable, as they appear correct just like in the example shown in figure 7.

V. Conclusions and future work

A method for analyzing range images, locating occlusions of large homogeneous areas and reconstructing these areas behind occluding objects has been presented. Some results obtained from the research have also been shown. From the work done so far, some remarks can be made. Partition by depth discontinuities and reconstruction did not cause any particular problems. The most difficult tasks are to partition the image based on fold edges and to make a correct choice of the regions to reconstruct. In both cases a major disturbance factor is the high level of noise encountered in range images; fold edge identification requires strong smoothing of the data and various devices to get around the consequent information loss. Noise is also a cause of error in region matching, as its influence on the parameters computed for small regions is overwhelming.

Decisions about region matching and the choice of areas to be reconstructed are currently based on low level information e.g. the location of a point with respect to the plane best fitting a region; the large number of odd cases encountered in range images would be better dealt with using a higher-level approach, such as recognition of observed objects. Progress in this direction may come from availability and use of knowledge on the range image content, though this is not straightforward. Other future work on this topic may be as follows:

1. A wavelet analysis of the visible area surrounding the reconstructed occlusion area may allow a more sensible reconstruction, taking into account structures like bumps that would otherwise be smoothed by depth interpolation.
2. The reconstruction could be made to work on more general shapes in addition to planes. Many walls in old buildings are cylindrical, and even more complicated shapes can be found. Of course matching parts of the shapes could cause higher complexity levels, and probably some knowledge about the environment should also be integrated. Preliminary results on this extension are presented in [17].
3. One could also develop a reliability gauge for the reconstructed pixel, as not all the pixels can be reconstructed with the same confidence.

Finally, one might argue that the solution to the occlusion problem is to simply acquire additional images. Even if it is possible to return the the identical scene, and to place the sensor in a location for scanning the missing area, recent results by Sanchiz [16] show that even simple scenes can require hundreds of images to obtain complete, high quality range data. Thus, it may be preferable to reconstruct small regions of missing data instead of attempting to observe them.

Acknowledgements

The authors wish to acknowledge SMART2 EC TMR network for funding this research.

References

[1] P. Besl, "Active, optical imaging sensors", Machine Vision and Applications, 1988, pp. 127-152.

[2] P. J. Besl, R. C. Jain, , "Segmentation through variable-order surface fitting", IEEE Trans. Pat. Anal. and Mach. Intel., vol. PAMI-10, No. 2, pp. 167-192, March 1988.

[3] B. Bhanu, L. A. Nuttall, "Clustering based recognition of occluded objects", in Proc. ICPR86, pp. 732-734, 1986.

[4] J. Bresenham, "Incremental line compaction", The Computer Journal, 25(1), pp. 116-120, 1982.

[5] CAMERA EC TMR network home page [Online]. Available: http://www.dai.ed.ac.uk/daidb/people/homes/rbf/CAMERA/camera.htm

[6] D. W. Eggert, A. W. Fitzgibbon, R. B. Fisher, "Simultaneous registration of multiple range views for use in reverse engineering of CAD models", Computer Vision and Image Understanding, 69(3), pp. 253-272, March 1998.

[7] R. B. Fisher, From Surfaces to Objects, J. Wiley & Sons, 1989.

[8] R. B. Fisher, A. W. Fitzgibbon, M. Waite, M. Orr, E. Trucco, "Recognition of complex 3-D objects from range data", in Proc. CIAP93, pp. 509-606, 1993.

[9] Y. K. Ham, R. H. Park, "3D object recognition in range images using hidden Markov models and neural networks", Pattern Recognition, 32(5), pp. 729-742, May 1999.

[10] R. L. Hoffman, "Object recognition from range images", Ph.D., Department of Computer Science, Michigan State University, 1986.

[11] R. Hoffmann, A. K. Jain, "Segmentation and classification of range images", IEEE Trans. Pat. Anal. and Mach. Intel., vol. PAMI-9, no. 5, pp. 608-620, Sept. 1987.

[12] A. Hoover, G. Jean-Baptiste, X. Jiang, P. J. Flynn, H. Bunke, D. Goldgof, K. Bowyer, D. Eggert, A. Fitzgibbon, R. Fisher, "An experimental comparison of range segmentation algorithms", IEEE Trans. Pat. Anal. and Mach. Intel., Vol 18(7), pp. 673–689, July 1996.

[13] S. L. Hurt, A. Rosenfeld, "Noise reduction in three-dimensional digital images", Pattern Recognition, 17(4), pp. 407-421, 1984.

[14] X.Y. Jiang, H. Bunke, U. Meier, "High-level feature based range image segmentation", Image and Vision Computing, vol. 18, No. 10, pp. 817-822, July 2000.

[15] P. G. Mulgahonkar, C. K. Cowan, J DeCurtins, "Understanding object configurations using range images", IEEE Trans. Pattern Anal. and Mach. Intel., 14(2), pp. 303-307, February 1992.

[16] J. M. Sanchiz, R. B. Fisher, "Environment recovery by range scanning with a next-best-view algorithm", Robotica, to be published.

[17] F. Stulp, F. Dell'Acqua, R. B. Fisher, "Reconstruction of surfaces behind occlusions in range images", Proc. 3rd Int. Conf. on 3-D Digital Imaging and Modeling (3DIM01),Montreal, Canada, June 2001.

[18] E. Trucco, R. B. Fisher, "Experiments in curvature-based segmentation of range data", IEEE Trans. Pat. Anal. and Mach. Intel., vol. PAMI-17, No. 2, pp. 177-182, February 1995.

[19] M. A. Wani, B. G. Batchelor, "Edge-region-based segmentation of range images", IEEE Trans. Pat. Anal. and Mach. Intel., vol. 16, no. 3, pp. 314-319, March 1994.

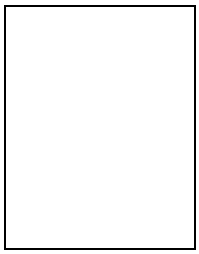

Fabio Dell'Acqua was born in Pavia on March 28th, 1971. In 1996 he got a first-class honour degree in Electronics Engineering at the University of Pavia. In 1999 he got a Ph.D. at the University of Pavia, investigating shape analysis techniques to process meteorological images. From September to December 1998 he was a visiting researcher at the Colorado State University in Fort Collins, Colorado, USA, to learn about processing of weather satellite images. In the first half of year 2000 he worked as a research associate at the Vision Laboratory of the University of Edinburgh in Edinburgh, Scotland, on analysis and reconstruction of range data (EU TMR - CAMERA). Since July 2000 he is a Post Doc at the university of Pavia. He is currently working on satellite image processing. His fields of interest include shape analysis, retrieval of images from archives, range data analysis, video compression (MPEG-4 and MPEG-7), neural networks and remote sensing.

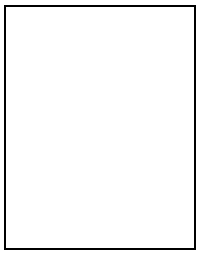

Robert B. Fisher received a B.S. with honors (Mathematics) from California Institute of Technology (1974) and a M.S. (Computer Science) from Stanford University (1978). He received his PhD from University of Edinburgh (1987), investigating computer vision in the Department of Artificial Intelligence. Dr. Fisher is a Reader in the Division of Informatics at the University of Edinburgh and is a member of the Institute of Perception, Action and Behaviour. His research covers topics in high level and 3D computer vision. He directs research projects investigating three dimensional model-based vision and automatic model acquisition of industrial objects and buildings. He teaches general and industrial vision courses for undergraduate, MSc and PhD level students.