

Environment Authentication through 3D Structural Analysis

Toby P. Breckon and Robert B. Fisher

Institute of Perception, Action and Behaviour
School of Informatics, University of Edinburgh
Edinburgh, UK

toby.breckon@ed.ac.uk / rbf@inf.ed.ac.uk
<http://www.ipab.inf.ed.ac.uk>

Abstract. We address the validation of the current sensed environmental state against a model known from earlier perception. Our surface based approach compares 3D range data of a built environment against an *a priori* scene model with the analysis of identified differences facilitating the hypothesis of causal structural changes within the scene. Experimental results show good success rates in identifying and analysing realistic structural changes introduced across example industrially themed scenes.

1 Introduction

Environment authentication is used to describe the validation of the current built environment against some *a priori* model in order to ascertain where changes may have occurred. Here we look at this technique with reference to building interiors where several applications, notably in the domain of mobile robotics and the nuclear/chemical safeguards industry, call for robust and generalised change identification and analysis methods [1]. In the latter case, the primary concern is to detect covert changes that may compromise safety or security [2].

Previous work in this area falls into two related but distinct categories - firstly the general topic of scene and object recognition and secondly the more specific area of scene change identification. The former, with specific relation to range data, has been well established through previous studies [3] of which [4] brought several of these techniques together to present a complete range-based 3D recognition system dealing with each stage of acquisition, segmentation and recognition. This work embodied a clear architecture for use in this field, primarily dividing the process into two tiers - the model independent tier involving the acquisition and segmentation of a 3D scene image and the model based reasoning tier handling issues of successful model to scene surface matching. Here we adopt this architecture for scene recognition and extend it by adding an additional post-process of structural change reasoning.

Work on the analysis of structural change within scenes is reasonably limited. Fillatreau et al [5] considered feature mis-placement with regard to validating

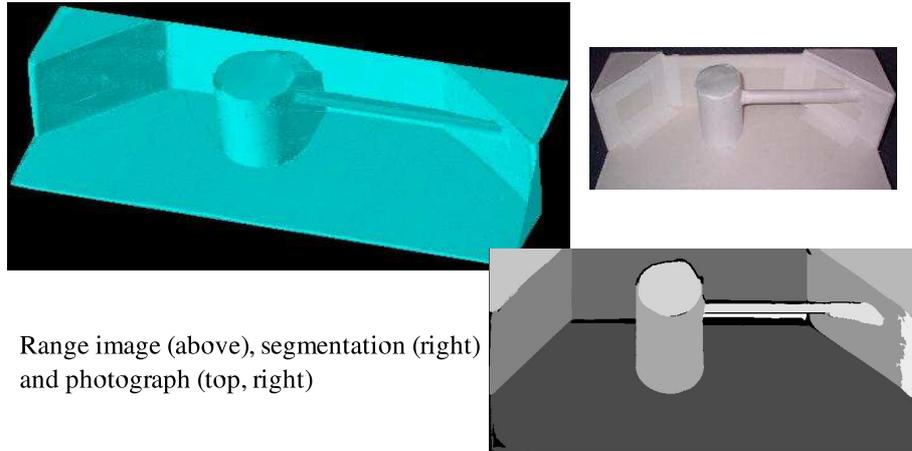


Fig. 1. Scene segmentation example

the correctness of an *a priori* model based on two stage tight and then relaxed constraint based matching. However, in this work [5] perfect segmentation is assumed and only localised feature movement is considered within the scene. Additionally, a significant body of work also exists within the remit of remote sensing (e.g. [6,7]) but this fails to address structural analysis within the localised environment concentrating instead on aerial or satellite image analysis. Similar work [8,9,10] has concentrated more on image differencing than on structural change detection, whilst work specifically with the remit of environment authentication [11] has shown the limitations of range occupancy grids in addressing problems in this domain.

Here we examine an alternative approach that targets full scene understanding through the use of classical surface matching techniques, augmented with additional structural reasoning, to provide a causal hypothesis for structural changes present in a possibly mis-segmented environment. Our results show that a surface based approach combined with later reasoning, based on spatial and geometric scene awareness, can provide a successful approach to 3D environment authentication.

2 Structural Analysis

Here structural analysis is based upon the comparison of a range image of the current environment against a known scene model. For this work we limit ourselves to the consideration of simple, industrially themed building interiors - where the geometric scene nature lends itself well to established model registration and comparison techniques [12].

Initially segmentation is performed using mean and Gaussian curvatures, with additional region growing and surface fitting techniques to provide a sur-

face map of the sensed environment (Fig. 1) together with a parameterised surface description - surface type¹, orientation, position and also radii where appropriate [13]. Due to the potential effects of noise on range imaging some mis-segmentation is to be expected and must be isolated from true cases of structural change.

For simplicity, our *a priori* scene model is represented using a basic scheme of generalised polyhedral, circular and cylindrical surfaces although a VRML or CAD model could similarly be employed.

The initial stages of processing follow the common invocation, matching and verification architecture of [4,12]. Our work here, however, focuses on a later stage of post-processing - *match analysis* of environment to model differences with a view of hypothesising causal structural changes within the scene.

Prior to detailing this aspect of our work in depth we briefly describe our classical approach in these earlier stages of processing. Firstly, invoked matching is used to produce a 'coarse match which is fast and inexpensive' based on the common position invariant surface attributes readily available from our representations. This lightweight match is then refined using a standard interpretation tree matching approach to provide a set of mutually consistent scene model to data surface matches. Here matching is considered through the consideration of unary and N-ary surface consistency - the matching of the individual and relative surface positions and orientation (within defined noise tolerances identified from ground truth environmental data) [3,12].

Once this matching process has identified a consistent set of surface matches, with a subset containing mutually non-parallel and independent surface orientations, the established SVD least squares fitting method of [14] is used to calculate the model to scene registration.

This registration is then used to both verify existing surface matches and find further matches in the scene based on a process of surface reprojection. Each scene surface is projected onto the model using the known registration and geometrically tested for correspondence against *a priori* model surfaces [15,12]. Once this verification process is complete all possible surface matches that can be found under normal strict matching conditions have been identified.

The next stage, an augmentation to the classical recognition architecture of [4,12] and extending the earlier work of [5], considers the analysis of the remaining unmatched surfaces by using relaxed matching to target full scene understanding through explanation of these occurrences - *match analysis*.

Match Analysis

The match analysis stage has two goals: 1) to isolate unmatched surface cases occurring due to occlusion and mis-segmentation; 2) to form a structural change hypothesis for remaining unmatched surfaces.

¹ {plane — cylinder — cone — sphere — general quadric} - however, only planar and cylindrical surfaces are fully considered in this work.

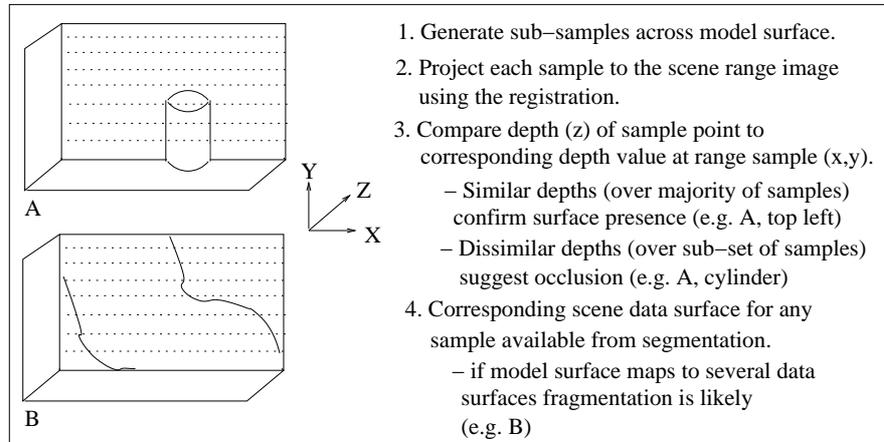


Fig. 2. Fragmentation / occlusion identification

Initially the remaining unmatched *a priori* model surfaces are screened to eliminate back facing (i.e. hidden) surfaces given the known model registration (back face culling, [16]). Next, a stage of occlusion and fragmentation analysis identifies surface matches missed by earlier processing because either they are a) partially occluded within the scene or b) they have been mis-segmented into one or more smaller **surface fragments**. Here these occurrences are detected by sub-sampling each unmatched model surface and projecting it onto the scene using the known model registration. By comparing the expected depth (model) and the measured depth (range image) at each sample, the true presence of these surfaces in the scene can be ascertained and instances of partial surface occlusion and fragmented surface presence identified (Fig. 2). This allows surfaces to be successfully matched to an occluded partner or a set of occurring surface fragments.

The final stage of this pre-classification analysis is to perform a ‘mop-up’ of remaining segmentation errors occurring within the scene data by identifying data fragments within the segmented scene. A **data fragment** is a surface considered to be too small to be of realistic significance and most likely to have been created erroneously due to the effects of noise on surface segmentation. Such surfaces are identified, based on size, and discarded from the remainder of the analysis.

Once these cases occurring due to mis-segmentation and occlusion have been identified and removed, the proper structural analysis can commence. The remaining unmatched scene and model surfaces are now classified as structural change based on four possible instances:

- **Movement:** when a surface match can be found by relaxing the matching constraints to consider only the position invariant surface attributes and

there exists further evidence, based on relative surface position, that the surface has moved within the scene.

- **Shape Change:** when a surface match can be found within the corresponding locale of the scene but differences in the shape attributes indicate the surface has changed shape within its original position.
- **Missing:** when no surface can be found to match an *a priori* model surface.
- **New:** when no model surface can be found to match an surface present in the scene environment.

In practise this classification uses the following top-down rule-set, which is evaluated for each remaining unmatched scene surface in terms of how it can be matched to a corresponding surface present in the model:

- *IF* match(size / shape attributes) *AND NOT* match(position)
 - **Movement**
- *ELSE IF* match(type, orientation and position)
 - **Shape Change**
- *ELSE*
 - **New**

Any remaining unmatched model surface is classified as **Missing**. All matching is performed based on identified noise tolerances which are derived empirically from measuring scene noise levels in ground truth environment to model matches.

In summary, our match analysis process operates sequentially as two components - firstly it eliminates possible causes of erroneous structural change detection by identifying occlusion and mis-segmentation instances, and secondly it classifies the remaining structural changes based on available evidence within the scene. Once this process is complete, every surface within the scene has been accounted for as either present, hidden, mis-segmented or forming part of a hypothesis explaining the structural change present in the scene. Overall our target of full scene understanding has been achieved.

3 Results

The match analysis process described above was tested over example structural change scenarios containing both single and multiple structural change occurrences (Table 1).

These were constructed using range scans of scale building interior models scanned with a 3D Scanners Reversa 25 laser range scanner² (e.g. Fig. 1). The type of changes introduced varied from subtle changes in surface positions and sizes to more complex and realistic changes such as alterations to plant within a scene (e.g. Fig. 3) and the detection of false walls / sealed doorways (e.g. Fig. 3 / Fig. 4). Figure 3 shows the successful detection of a false back wall (A), occurring

² x/y resolution: 0.4mm, depth accuracy (z): 50 microns

	Scenario	% Success	Surface Examples Tested
G	Ground truth match	94%	104
S	Single changes within the scene	87.5%	56
M	Multiple changes within the scene	87%	46
F	Surface fragmentation within the scene	88%	17

Table 1. Experimental Results

in the scene forward of its expected model position, as well as movement (D) and change (B) in the central ‘water tank’ configuration. Additional surface fragments (E / F), hidden surfaces (G / H) and a missing support cylinder (C) are similarly identified. Figure 4 shows the successful detection of a sealed doorway in the central rear wall of the scene (A) and again successful surface fragment detection (B/C).

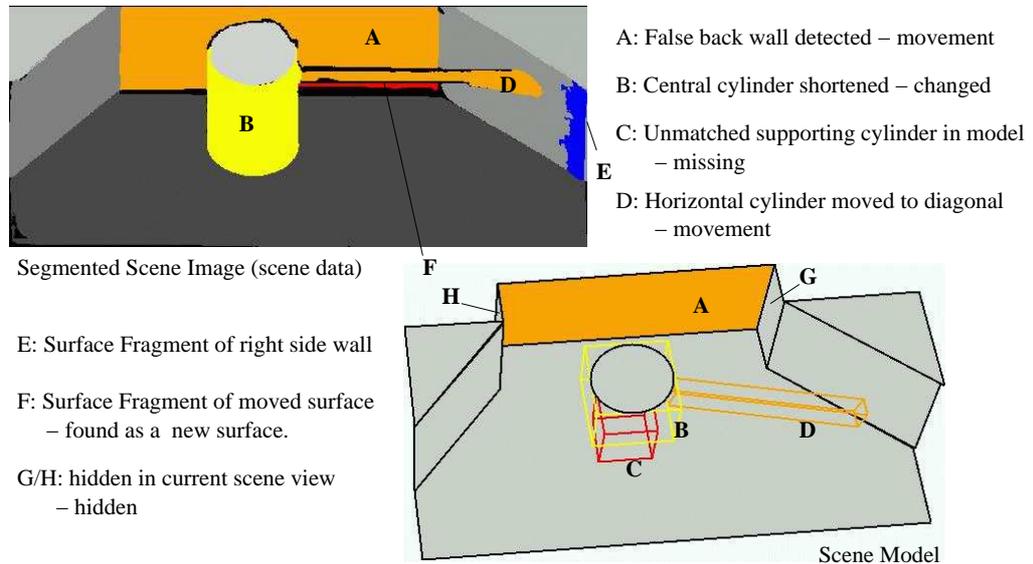


Fig. 3. Successful false wall, plant changes and surface fragment detection.

Overall the system was successful in detecting 87% of the 102 structural changes (cases S and M from Table 1) introduced to the test scenes. In the remaining cases a combination of poor segmentation and surface fitting, due to scene noise in the range capture, reduced sensitivity. Lesser sensitivity is equally apparent when surfaces move subtly within their defining plane and when surfaces are heavily occluded within the scene. This is due to the limited

availability of structural change evidence outwith the tolerances of regular scene noise. Additionally, the process is susceptible to any errors present in achieving a ground truth match and to identifying movement cases in mis-segmented surfaces (Fig. 3, item F).

Segmentation related errors were also successfully detected in testing with all data fragments and 88% of surface fragments being correctly identified (case F from Table 1, e.g. Fig. 3, item E). In all cases where surface fragments were missed a high level of scene noise caused poor segmentation making successful surface matching difficult. Further advances in noise tolerant segmentation and 3D data acquisition may help to counter this issue.

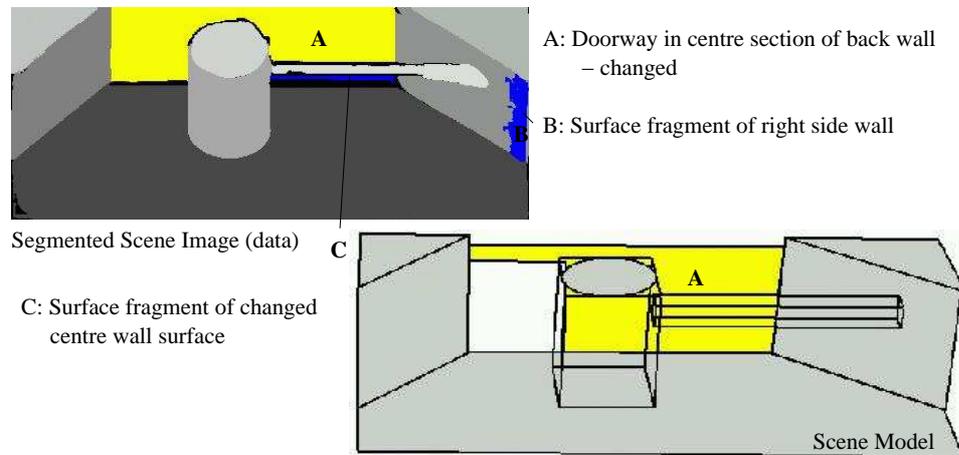


Fig. 4. Detection of sealed doorway and surface fragments.

4 Conclusions

We have presented a method for environment authentication and structural change analysis through full scene understanding of a segmented environment range image. Here we overcome the limitations of earlier work [11,8,5] by identifying individual surface changes, handling mis-segmentation and occlusion based issues and performing analysis to form an explanative hypothesis for identified structural changes.

Although the results quoted show a significant level of success for these techniques, further work is still required in a number of areas. Notably, improvements in range image segmentation and surface extraction together with those in occluded surface reconstruction [17] offer possibilities for improving accuracy and reducing the reliance on current mis-segmentation and occlusion handling techniques. Similarly, improvements in match analysis techniques to consider more advanced best fit based hypothesis construction, the use of further scene analysis

evidence metrics and also the use of probability based classification may be beneficial in future work. Additionally, the extension of this or similar techniques, to consider the matching of meshes representing curved surfaces and the consideration of higher-level ‘changed structures’ as consistent groupings of underlying surface changes may be of future interest in advancing work in this area.

Acknowledgements This work was partially supported by studentship funding from the EPSRC and CASE sponsorship by QinetiQ PLC (<http://www.qinetiq.com>).

References

1. Latombe, J.: Robot Motion Planning. Kluwer (1991)
2. Government, H.M.: Confidence, security and verification - the challenge of global nuclear weapons arms control. Technical report, Atomic Weapons Establishment, UK (2000)
3. Grimson, W.: Object Recognition by Computer: The Role of Geometric Constraints. MIT Press (1991)
4. Fisher, R., Fitzgibbon, A., Waite, M., Trucco, E., Orr, M.: Recognition of complex 3D objects from range data. In Impedovo, S., ed.: Proc. 7th Int. Conf. on Image Analysis and Processing, Monopoli, Bari, Italy (1993) 509–606
5. Fillatreau, P., Fisher, R.: Non-polyhedral landmark recognition using 3D depth images and partial correct models. In: Proceedings 3rd Int. Symp. Intelligent Robotic Systems. (1995) 271–278
6. Ridd, M., Liu, J.: A comparison of four algorithms for change detection in an urban environment. Remote Sensing of Environment **63** (1998) 95–100
7. Huertas, A., Nevatia, R.: Detecting changes in aerial views of man-made structures. Image and Vision Computing **18** (2000) 583–596
8. Benjanin, M., Huertas, A., Medioni, G., Nevatia, R.: Model validation for change detection. In: DARPA Image Understanding Workshop, Monterrey, CA, USA. (1994)
9. Rosin, P.: Thresholding for change detection. In: Int. Conf. on Computer Vision. (1998) 274–279
10. Murakami, H., Nakagawa, K., Hasegawa, H., Shibata, T., Iwanami, E.: Change detection of buildings using an airborne laser scanner. ISPRS J. of Photogrammetry and Remote Sensing **54** (1999) 148–152
11. Campos, G., Goncalves, J.: Environment authentication using range occupancy grids. JRC Technical Note I93.130, EC JRC, Institute for Systems Engineering and Informatics, Ispra, Italy (1993)
12. Faugeras, O.: Three-dimensional computer vision: a geometric viewpoint. MIT Press (1993)
13. Fisher, R., Fitzgibbon, A., Eggert, D.: Extracting surface patches from complete range descriptions. In: Proc. Int. Conf. on Recent Advances in 3D Digital Imaging and Modelling, Ottawa, Canada (1997)
14. Arun, K., Huang, T., Blostein, S.: Least-squares fitting of two 3D point sets. IEEE Transactions on Pattern Analysis and Machine Intelligence **9** (1987) 698–700
15. Haines, E.: Point in polygon strategies. In Heckbert, P., ed.: Graphics Gems IV, Academic Press (1994) 24–46
16. Foley, J., van Dam, A., Feiner, S., Hughes, J.: Computer Graphics: Principles and Practice. Addison-Wesley (1996)
17. Castellani, U., Livatino, S., Fisher, R.: Improving environment modelling by edge occlusion surface completion. In: Proc. Int. Symp. on 3D Data Processing Visualisation and Transmission, Padova, Italy (2002)