New Techniques for Automated Architectural Reconstruction from Photographs

Tomas Werner and Andrew Zisserman

Robotics Research Group
Department of Engineering Science
University of Oxford
Oxford, OX1 3PJ
{werner,az}@robots.ox.ac.uk

Abstract. We investigate a strategy for reconstructing of buildings from multiple (uncalibrated) images. In a similar manner to the Facade approach we first generate a coarse piecewise planar model of the principal scene planes and their delineations, and then use these facets to guide the search for indentations and protrusions such as windows and doors. However, unlike the Facade approach which involves manual selection and alignment of the geometric primitives, the strategy here is fully automatic.

There are several points of novelty: first we demonstrate that the use of quite generic models together with particular scene constraints (the availability of several principal directions) is sufficiently powerful to enable successful reconstruction of the targeted scenes. Second, we develop and refine a technique for piecewise planar model fitting involving sweeping polygonal primitives, and assess the performance of this technique. Third, lines at infinity are constructed from image correspondences and used to sweep planes in the principal directions.

The strategy is illustrated on several image triplets of College buildings. It is demonstrated that convincing texture mapped models are generated which include the main walls and roofs, together with inset windows and also protruding (dormer) roof windows.

1 Introduction

The objective of this work is the automated construction of 3D texture mapped models of buildings from a number of close-range photographs of the scene. Much research into automation in the photogrammetry community has concentrated on reconstruction from aerial, i.e. long-range, views (e.g. [3,5,9,15,19, 20]). For close-range photographs there are several manual systems available (e.g. Facade [27], Canoma [1], PhotoModeller [2]) which may be used to produce excellent piecewise polyhedral and piecewise quadric scene models of buildings.

Automated close-range reconstruction of buildings [11,12,21,26], is less thoroughly explored than long-range. Two recent papers represent well the spectrum of approaches currently available. Dick *et al* [12] cast the problem as model based recognition, and attempt to fit very strong models using a Bayesian framework. The models are highly tuned to the scene (e.g. a gothic window for a chapel)

A. Heyden et al. (Eds.): ECCV 2002, LNCS 2351, pp. 541–555, 2002.

and the strong priors (e.g. on height and width) considerably narrow down the search. The method recognises these models in a single image, and other images may then be used to verify these recognition hypotheses. At the other end of the spectrum Tao *et al* [26] cast the problem as that of fitting piecewise planar patches over multiple views. The patches are determined by (over-) segmenting a reference image using colour, and subsequently larger regions are assembled by merging 3D patches. In contrast to [12] the method cannot proceed if only a single view is available as it is completely rooted in multiple view stereo. Both methods produce excellent results.

In this paper we explore a strategy that is located in the 'middle' of this spectrum: we fit quite generic models (planes, polyhedra). These are weaker than the scene specific models of [12] but stronger than the planar patches employed in [26]. Our approach is inspired by that of Facade [27] and proceeds in two stages: first, a coarse piecewise planar model of the principal scene planes and their delineations is generated; and second, this coarse model is used to guide the search for fitting more refined polyhedral models of the indentations and protrusions from the coarse model, such as windows and doors.

The method is targeted on architectural scenes which typically contain planes orientated in three dominant directions which are perpendicular to each other, for example the vertical sides of a building and the horizontal ground plane. The strategy assumes that the scene contains three such principal directions and that the images contain sufficient information to obtain the vanishing points of these directions.

At the technological level we considerably extend the one parameter 'plane-sweep' method that has been used by several previous researchers for plane fitting [4,8,11,16]. The extensions are of two types: previously the method was employed with calibrated cameras with known external orientation. Here we are in an uncalibrated setting, and the directions for the sweep are computed from vanishing points in the images. The second extension is to sweep not only planes, but also polyhedral models.

2 Preliminaries

In this section we describe the stages of camera computation, vanishing point detection, image point and line matching, and the estimation of the 3D counterparts of the points and lines. This is the necessary starting point in preparation for the 'clothing' of the scene by planar primitives. These stages do contain novel elements, but they are not the main issue of this paper so will only be described briefly in the following paragraphs. They will be illustrated for the image triplet shown in figure 1.

Projective Reconstruction: Interest points are matched between the image triplet and a trifocal tensor estimated using standard robust methods (see [17] for full details and citations). The result is a camera corresponding to each image and a set of 3D points, defined up to an unknown projective transformation of

(a) (b) (c)

Fig. 1. Three images of a Merton College court acquired with a hand held low cost Olympus digital camera. The image size is 1024×768 pixels.

3-space. The projective reconstruction computed from the images of figure 1 has 684 3D points, which after bundle adjustment achieve a RMS reprojection accuracy of 0.14 pixels, with a maximum error of 1.01 pixels.

Vanishing points: Much previous work has been devoted to automatic vanishing point detection (e.g. [6,7,10,18,22,23,25,28]), and we employ a standard method where lines are computed independently in each image, and the vanishing points corresponding to the three principal directions estimated using a RANSAC method. The typical line sets used to compute the vanishing points are shown in figure 2. The vanishing points are then matched across the images by a simple combinatorial algorithm. The principal direction (which is a point in 3D, the pre-image of a vanishing point) is estimated by minimising the reprojection residuals from the lines supporting each vanishing point.

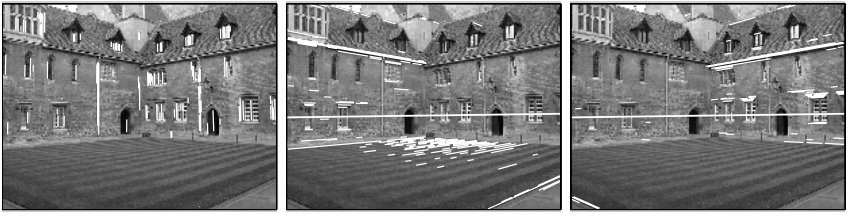

(a) (b) (c)

Fig. 2. Vanishing point computation. (a), (b) and (c) shows the lines (in white) supporting the vanishing points corresponding to the three principal directions for the first image of figure 1. The thick line in (b) and (c) is the horizon computed as the line through the vanishing points.

These principal directions (which are points at infinity) then play an important role throughout the reconstruction strategy. For example, in determining a vertical facade plane by a RANSAC search on 3D point triplets, the directions

provide two of the points with only one point required to be sampled from the matched scene points. This is described in section 3.

Metric rectification: The approach used here to upgrade the projective reconstruction to metric is based on two metric constraints: the three principal and orthogonal scene directions; and that the cameras have square pixels. The method proceeds in three steps. First, the three principal directions (points at infinity) are computed from their images (the vanishing points). This determines the plane at infinity as the plane defined by the three directions, and consequently the projective reconstruction is upgraded to affine. Second, the principal directions are forced to be mutually orthogonal. This determines a reconstruction which differs from metric only by a scaling in each of the principal directions. In the last step, these scalings are determined (up to an overall scale) using the constraint that the camera pixels are square by linearly minimizing an algebraic residual.

This linear, non-iterative algorithm yields a very good initial estimate of the metric cameras, i.e. both the internal and exterior orientation. For example the computed internal parameters for the first image are: principal point (580.5, 349.0), aspect ratio 1.00055, angle between image axes 89.8944°, and focal length 1085.4 pixels. Their further improvement is possible by a bundle adjustment constrained by orthogonality and square pixels assumptions. The reconstructed 3D point cloud after the metric rectification is shown in figure 3.

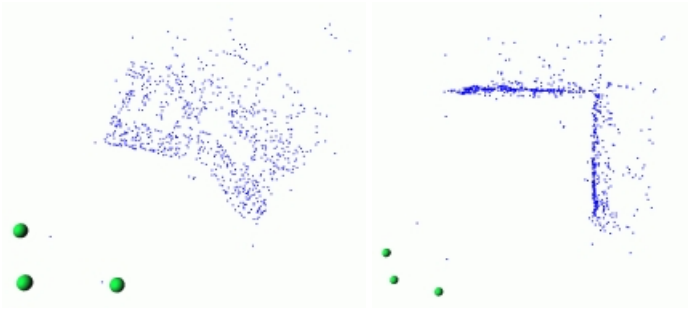

Fig. 3. Two views of the cloud of 3D points after metric rectification. The spheres depict the camera centers.

Line matching: We use a local implementation of the line matching method of Schmid and Zisserman [24]. Line segments are detected independently in each image, and then matched across the triplet using a combination of photometric constraints and trifocal geometry. In total 434 line segments are matched. Special attention is paid to horizontal lines as these tend to lie in epipolar planes for

typical acquisition geometries. The resulting ambiguity is resolved by a search to register the photometric neighbourhood – a 'line sweep'.

We then carry out another stage of processing where the lines are classified into disjoint direction sets and re-estimated with this additional constraint. This is crucial for (i) accuracy of the estimated 3D lines, and (ii) organizing the subsequent search for higher level building primitives (planes, polygons). In more detail, from the intersection point of the line with the plane at infinity, the lines are labeled (using χ^2 statistics on reprojection residuals) as belonging to one of seven classes: principal direction (3), principal line (3), other (1). For example if a line is vertical, then it intersects the plane at infinity at the vertical direction. When it is re-estimated under this (hard) constraint, the minimization of reprojection error is only over the two remaining parameters of the line. Figure 4 shows the resulting matched lines in 3D.

(a) (b)

Fig. 4. Reconstructed 3D lines, shown reprojected onto the first image (a) and in 3D (b). Of 193 matched lines, there are 7 mismatches (mostly due to repeated structure on the lawn). Of the rest, 168 lines go through a principal direction (e.g., vertical), 19 go through a principal line (e.g., are parallel to a wall), and none is unconstrained.

3 Coarse Model Fitting

The objective of this section is to automatically construct a coarse polyhedral model of the scene, given the cameras, 3D points, directions and lines, whose computation was described in the previous section. Numerous methods have been given for plane and polyhedral fitting and we employ these with the addition of direction constraints.

In overview we organize the search in the following stages: first we search for planes intersecting two principal directions (typically walls). The intersections of these planes give a partial delineation, and visibility is then used to determine

which parts of the delineation are included. Second, we search for planes incident with the selected wall planes and intersecting one principal direction (roofs). The vertical walls are consequently completely delineated: above by the roof plane, below by the ground plane, and laterally by other planes or the image boundary.

To give an example of a plane search, suppose we are searching for the vertical wall. Two types of methods are employed. One is the standard RANSAC [13] robust estimation. The novelty here is that samples are not simply based on three random points, but on one of the following particular samples: one point and a 3D line of the appropriate class (in this case not oblique); one point and the appropriate line at infinity (the one corresponding to vertical planes); two points and the appropriate point at infinity (e.g. the vertical direction). This ensures that only planes of the required type are estimated, and has a considerable advantage in reducing complexity – e.g. one finite 3D point match is sufficient to instantiate a plane. The second type of method is plane sweeping which is described in more detail below. Both types of method are attempted and the result of the most successful is used. The other planes (e.g. the roof planes) are searched for in a similar manner.

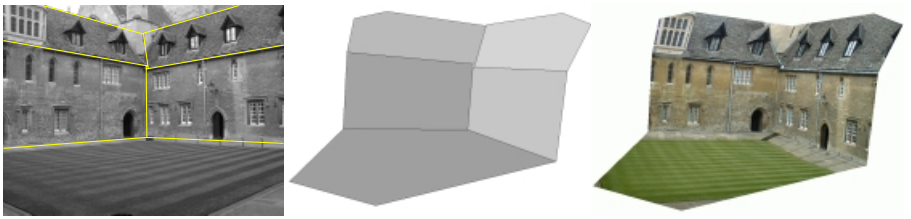

(a) (b) (c)

Fig. 5. Coarse model: wire frame projected to image 1 (a), shaded (b) and texture mapped (c) 3D model.

In most cases to date finding the principal vertical planes is straightforward. The ground plane often contains fewer point and line matches, and is found by plane sweeping. Once the vertical and ground planes are known, the search for roof planes is simpler because much of the erroneous data can now be ignored. The coarse model is shown in figure 5.

3.1 Plane Sweeping

The objective here is to determine scene planes by a one parameter search (the sweep) of a virtual plane. The support for the virtual plane as it is swept is measured using the planar homography induced by the plane to define a point-to-point mapping between the images, and aggregating cross-correlation scores directly from the images. The novelty of the application here is that we define the sweep about a line at infinity. A line determines a plane up to one parameter,

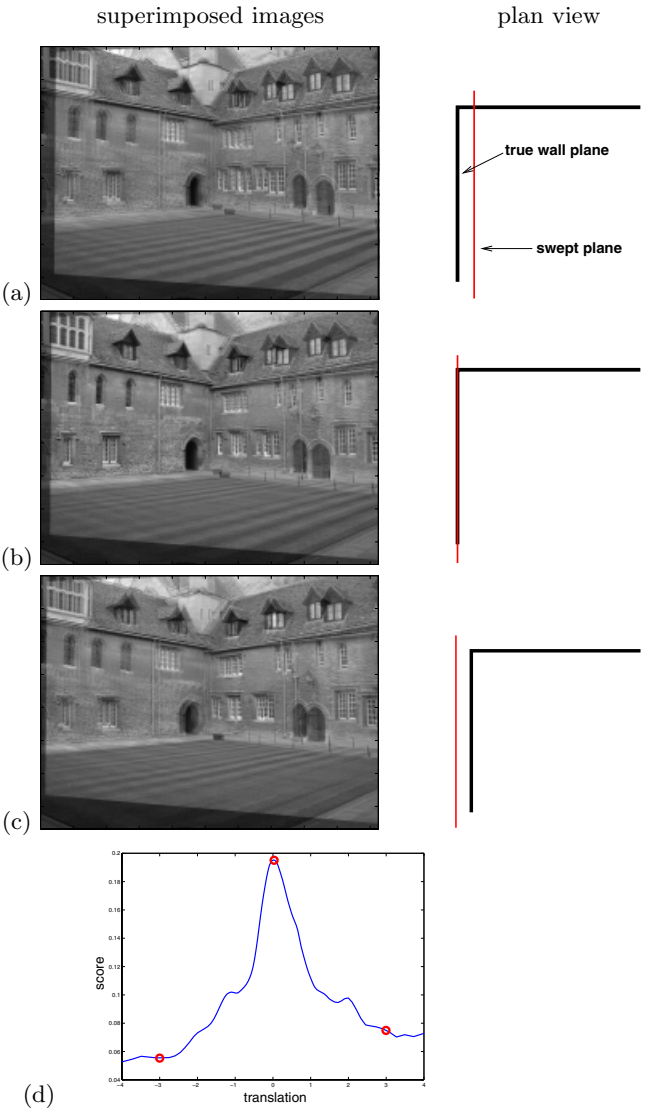

Fig. 6. Plane sweep under translations along a principal scene direction. (a,b,c) show images in figure 1a and 1b superimposed by a homography map corresponding to a translating a virtual scene plane. This scene plane is parallel with the left wall. The right-hand figures illustrates position of the wall and swept plane from a plan view. (d) shows a plot of the score function against translation. The circles correspond to the translations in (a,b,c), respectively. The middle translation is the one which best registers the planes, and this is visible in (b) where the plane of interest is most "focussed".

and in this case the parameter defines a translation of the plane since planes through a line at infinity are parallel.

For example, to search for the ground plane, the horizontal line at infinity is used, and the resulting sweep is in the vertical direction across the scene. When the swept plane corresponds to the actual ground plane, then image regions mapped by the homography will be highly correlated in the region corresponding to the ground. By determining the swept plane which produces the highest correlation (measured by an appropriate robust score function described below) the true position of the ground plane is determined. The sweeping is illustrated for one of the principal directions in figure 6.

Plane sweeping is able to determine scene planes even in the absence of any point or line correspondences for that plane (as are required in RANSAC fitting).

The cross-correlation is measured as follows: first image points at significant image gradients are determined (e.g., [14]); second, a sample of these points are selected as the key points at which to measure the cross-correlation. Only these points are used in order to improve the signal to noise ratio by discarding uninformative points in homogeneous image regions. Approximately 10% of image pixels are selected as the keypoints. At each of the key points the cross-correlation is

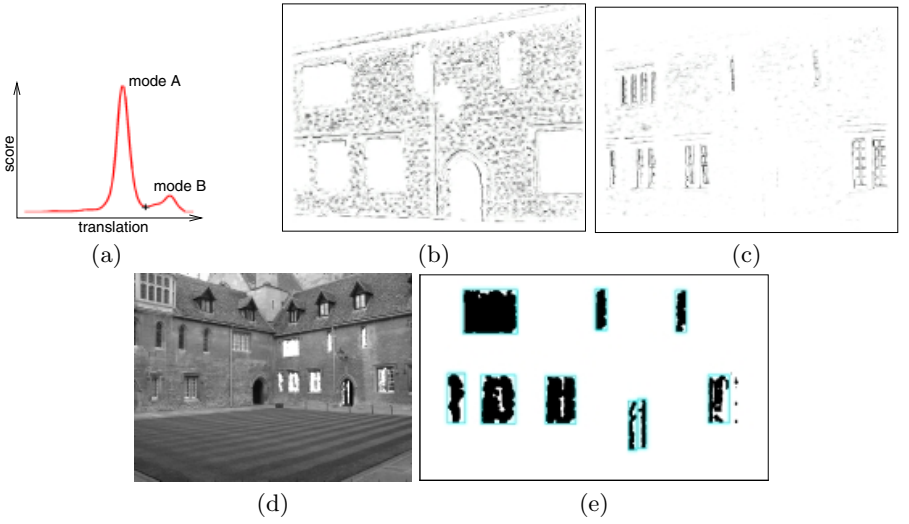

Fig. 7. (a) Aggregated similarity score in the neighbourhood of the wall plane computed over three views. Mode A corresponds to the wall plane, and B to the plane of the windows. The magnitude of the similarity scores for individual key points are represented by pixel darkness in (b) for the swept plane at the position of mode A, and (c) for the position of mode B. The key points belonging mostly to the wall plane score highly (are dark) in (b) and those belonging to the window plane score highly in (c). Points belonging to windows, (d), are obtained by thresholding depths at the value denoted by the cross in plot (a). (e) Rectangles robustly fitted to clustered points shown on the rectified facade.

computed between the point's neighbourhood and the corresponding neighbourhood in the other image mapped by the homography induced by the sweeping plane. A neigbourhood size of 7×7 pixels is used here. The cross-correlation is computed between all pairs of images, and the result averaged. A score function consists of the cross-correlations averaged over all points as a function of plane translation. Typical score functions are shown in figure 6.

Performance assessment. In practice plane sweeping is a very powerful fitting method, and we may ask why it is so successful. There are two main reasons: (i) it correctly models and corrects for distortion of the correlation region between the image, in a similar manner to wide baseline stereo matching (*cf.* stereo rectification which doesn't correct) – this is the reason that the matching of individual key points is so precise (see figure 7b and c), and (ii) the scoring function aggregates over a large number of points in order to estimate depth (*cf.* traditional stereo where the depth of each point is estimated independently).

4 Refinement

In this section we refine the coarse polyhedral model computed in the previous section by attempting to fit two types of polyhedral model to account for deviations from the plane. The models are a rectangular block (this fits to doors and windows, described in section 4.1), and a wedge block (this fits to dormer windows (in roofs) and wall protrusions, described in section 4.2).

These models are quite generic, and in contrast to the approach of [12] there is no use of strong priors to define (for example) their aspect ratio. Fitting is initiated in regions on either side of the fitted plane where there is evidence of a perturbation from the plane. All models are fitted by variations on sweeping as described below.

4.1 Fitting Rectangular Block Models

The idea is to determine regions of the plane which do not coincide with the coarsely fitted scene plane, and then to model these regions as rectangles aligned with the principal scene directions. Note, each region is modelled independently and from depth information alone.

Points which lie behind the fitted plane (indentations) are determined by thresholding depths of individual keypoints. The threshold value is obtained from the score function recomputed for the image region corresponding to the current facade, see figure 7a. Two modes are clearly discernable in the function – one corresponding to the coarse facade plane, and the other to the window plane.

The keypoints labeled by thresholding as being behind the wall are shown in figure 7d. Contiguous regions (corresponding to each of the windows and doors) are then computed by robustly clustering these points. The fitted plane has essentially simplified this task to that of clustering a set of pixels in a 2D

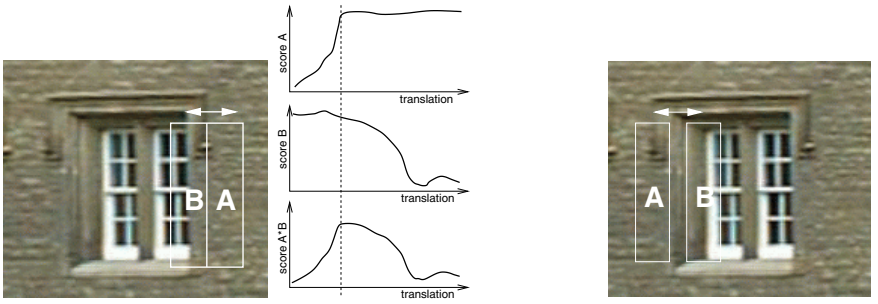

(a) (b) (c)

Fig. 8. Refining the position of the vertical window boundaries. (a) Search regions consisting of two adjacent rectangles which are translated horizontally. In rectangle A the similarity score is based on the wall plane. In rectangle B the window plane is used. (b) plots the two scores and their product, which peaks at the actual window boundary location (the dashed line). (c) For non-occluding window edges (the left one in this case), a gap is inserted between the rectangles. The width of the gap is the distance between the wall and window planes projected to the image.

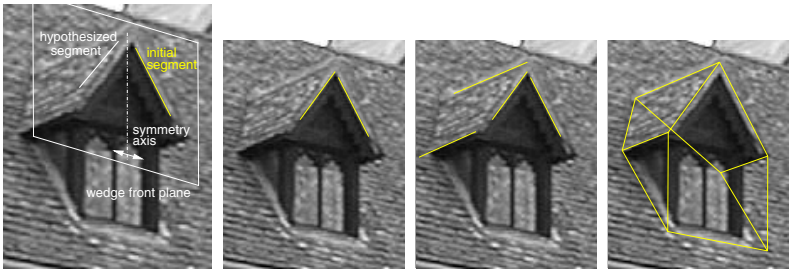

(a) (b) (c) (d)

Fig. 9. Fitting a wedge model (to reconstruct a dormer window). (a) Fitting is initiated if a single oblique 3D line segment (depicted in yellow) is detected. A symmetry related line segment is determined by searching for intensity gradient maximum, the gradient is summed along the hypothesized line segment. (b) The detected symmetric segment. (c) The top and lower line segments are determined by a similar search. (d) The model is then completed in a straightforward manner, using the known roof plane.

image as the analysis can be carried out on a rectified version of the image (where the principal scene directions are orthogonal in the image). Standard image processing methods, such as simple operations of binary morphology, are used. Rectangular boxes are then fitted to the clusters. The resulting windows boundaries are shown in figure 7e.

The window boundaries are then refined using a further correlation based search, but now concentrated in the vicinity of the putative window boundary. For each boundary a one-dimensional search is carried out perpendicular to that boundary. For example for a vertical boundary the search is horizontal. Two score functions are computed, one based on the homography induced by the

wall plane, the other based on the homography induced by the window plane, as illustrated in figure 8. For the wall plane the score is high when the pixels in the rectangle belong to the wall and small otherwise. Conversely, for the window plane homography the similarity score is high for window pixels, and low for wall ones. The product of these two scores peaks at the actual window boundary, as shown in the figure. In essence this is another sweep-like fitting of a model.

4.2 Fitting Wedge Models

Wedge shaped protusions are frequently found in architecture, for example they might form part of the roof or a bay windows. We illustrate here how a wedge model is instantiated and fitted automatically in order to reconstruct dormer windows in the vicinity of a roof plane.

The method is demonstrated in figure 9. Initially, all reconstructed scene lines that intersect exactly one principal line are considered as belonging to a potential wedge shape. The principal line of the segment defines the wedge front plane. If the other wedge edges have been matched and are available in 3D, finishing the wedge is straightforward. However, usually these lines are not available – for example in figure 1 many dormer edges are not even detected by the Canny edge detector due to the complicated roof texture. In this case a search for a line segment symmetrical to the initial segment is performed. The search is parametrized by a single parameter (the horizontal position of the symmetry axis) and the true location maximizes the intensity gradient along the segment. In a similar manner, the remaining wedge edges are localized. All dormer windows in figure 1 are successfully reconstructed.

4.3 Results

The facade model updated with models found in sections 4.1 and 4.2, and texture mapped from the appropriate images is shown in figure 10.

The results of applying the same algorithms (coarse then refinements) to an another image triplet are shown in figure 11. In this case, a fitting algorithm similar to one described in Section 4.2 automatically finds the protrusion in the left wall. The algorithm works on the same principle, but uses a different protrusion model. It is instantiated by the horizontal oblique edge at the top of the protrusion. The dimensions of the protrusion are then found by sweeping hypothesized edges in directions determined by the initial edge and the principal directions. High gradient along the swept edge is used as evidence for the true edge. Again, the classification of principal directions is crucial here.

5 Discussion

In this work we have systematically explored a particular strategy for scene reconstruction.

Clearly a strategy of this type is not applicable to non-urban scenes (trees, mountains) or cluttered interiors (rooms with tables, table legs, piles of books).

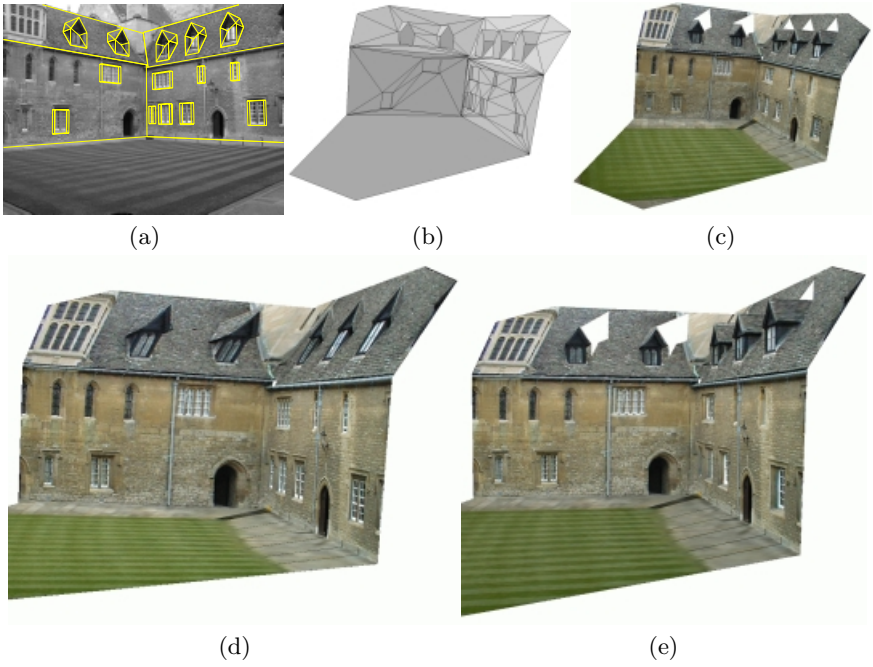

Fig. 10. The automatically computed 3D model augmented with indented windows in the vertical planes, and dormer windows in the roof. (a) wire frame projected onto the image, (b) shaded wireframe, (c) textured mapped 3D model. Compare the same view of the coarse model (d) and the augmented final model (e) – explicit modelling of dormer windows and indented windows on the walls is clearly an improvement. Note, in textured models the occluded parts of the roof and wall planes (e.g. behind the dormer windows) are removed and appear as white.

A more interesting question is where it will fail in its targetted domain of architectural scenes. By and large, and for obvious reasons, many buildings have many horizontal and vertical parts (bricks, edges, windows, doors) so that a vertical direction and line at infinity for horizontal plane may be estimated from images. The weakest point of the current strategy is that two orthogonal horizontal directions are not always available in the scene, e.g., on an octagonal building, or a circular church. The lack of these directions affects the strategy at several points. Fortunately, alternatives are available and these are currently being incorporated so that the strategy is not so limited. For example, when plane sweeping for perturbations the direction of the sweep can simply be chosen as perpendicular to the planes of the coarse geometric model.

One of the important contributions of this work is the development of fitting geometric models by sweeping, and three new techniques are described in the previous sections, namely sweeping scene planes about a line at infinity (used for obtaining or refining positions of the main walls in section 3.1), correlation

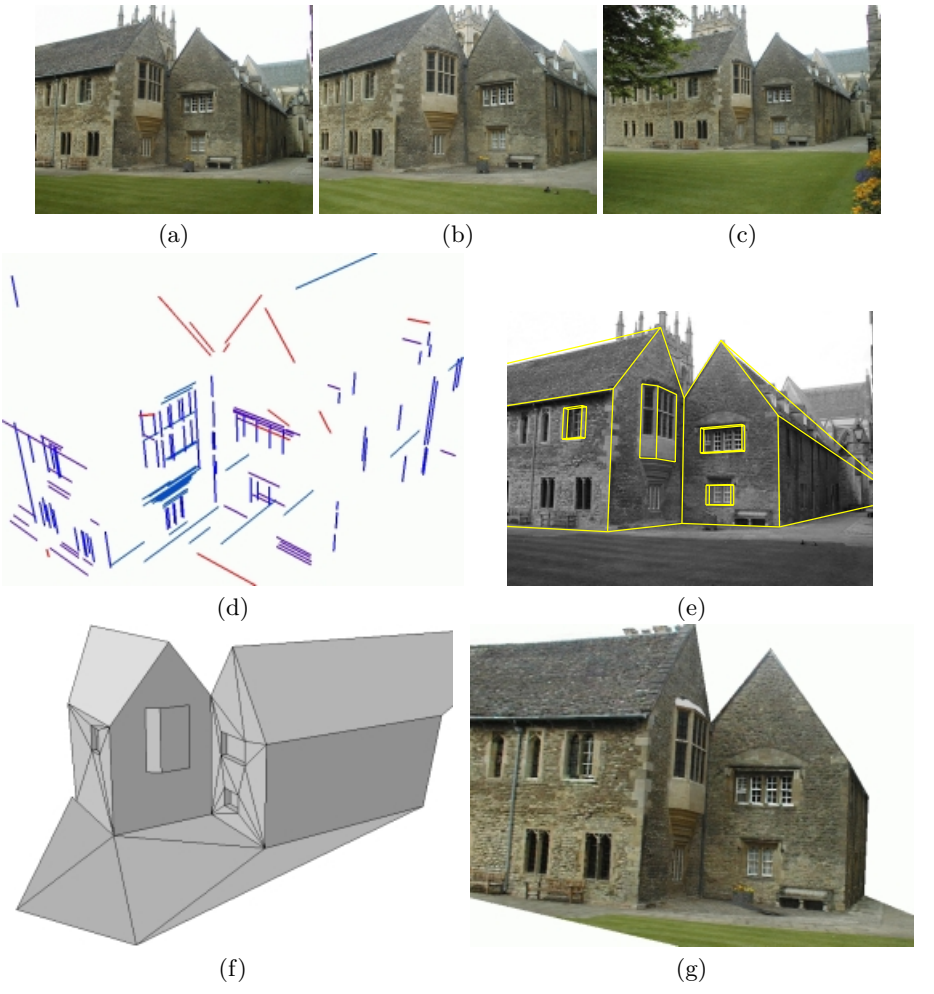

Fig. 11. (a,b,c) Input images of another court of Merton College, Oxford, acquired with the same camera as figure 1. Reconstruction results: (d) reconstructed and classified line segments, (e) model wire frame projected onto image 1, (f,g) final model with indented windows and protrusion reconstructed.

based search for building edges using translating rectangles and inter image homographies (used for refining window boundaries in section 4.1), and search for local gradient maxima along translating line segments (used for dormer window reconstruction in section 4.2). They have much in common: (i) A single unknown parameter, making the search fast and simple. (ii) Evidence is aggregated over relatively large region of interest, enabling features to be found that would not be detected by local operators such as Canny edge detectors or correlation based stereo on individual pixels. (iii) Results of the previous reconstruction stages

allow sweep algorithms to focus only on relevant image parts, achieving a high signal to noise ratio.

Acknowledgements. We are very grateful to Frederick Schaffalitzky for computing the projective reconstructions for the triplets. Funding was provided by a Marie Curie fellowship and EC project Vibes.

References

1. `http://www.canoma.org`.
2. `http://www.photomodeler.org`.
3. C. Baillard, C. Schmid, A. Zisserman, and A. Fitzgibbon. Automatic line matching and 3D reconstruction of buildings from multiple views. In *ISPRS Conference on Automatic Extraction of GIS Objects from Digital Imagery, IAPRS Vol.32, Part 3-2W5*, pages 69–80, September 1999.
4. C. Baillard and A. Zisserman. Automatic reconstruction of piecewise planar models from multiple views. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 559–565, June 1999.
5. F. Bignone, O. Henricsson, P. Fua, and M. Stricker. Automatic extraction of generic house roofs from high resolution aerial imagery. In *Proc. 4th European Conference on Computer Vision, Cambridge*, pages 85–96, 1996.
6. B. Brillault-O'Mahony. New method for vanishing point detection. *CVGIP: Image Understanding*, 54(2):289–300, 1991.
7. C. Coelho, M. Straforini, and M. Campani. Using geometrical rules and a priori knowledge for the understanding of indoor scenes. In *Proc. British Machine Vision Conference*, pages 229–234, 1990.
8. R. T. Collins. A space-sweep approach to true multi-image matching. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 358–363, 1996.
9. R. T. Collins, C.O Jaynes, , Y-Q Cheng, X. Wang, F. Stolle, E. M. Riseman, and A. R. Hanson. The ascender system: Automated site modeling from multiple images. *Computer Vision and Image Understanding*, 72(2):143–162, 1998.
10. R. T. Collins and R. S. Weiss. Vanishing point calculation as a statistical inference on the unit sphere. In *Proc. 3rd International Conference on Computer Vision, Osaka*, pages 400–403, December 1990.
11. S. Coorg and S. Teller. Extracting textured vertical facades from controlled close-range imagery. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition, Fort Collins, Colorado*, pages 625–632, 1999.
12. A. R. Dick, P. H. S. Torr, Ruffle S. J., and R. Cipolla. Combining single view recognition and multiple view stereo for architectural scenes. In *Int. Conf. Computer Vision*, pages 268–280. IEEE Computer Society, 2001.
13. M. A. Fischler and R. C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Comm. Assoc. Comp. Mach.*, 24(6):381–395, 1981.
14. W. Förstner and E. Gülch. A fast operator for detection and precise location of distinct points, corners and center of circular features. In *Proc. of ISPRS Intercommission Conference on Fast Processing of Photogrammetric Data, Interlaken, Switzerland*, pages 281–305, June 2-4 1987.

15. M. Fradkin, M. Roux, and H. Maître. Building detection from multiple views. In *ISPRS Conference on Automatic Extraction of GIS Objects from Digital Imagery*, September 1999.
16. A.W. Gruen. Adaptative least squares correlation: a powerful image matching technique. *S. Afr. Journal of Photogrammetry, Remote Sensing and Cartography*, 3(14):175–187, 1985.
17. R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521623049, 2000.
18. G. F. McLean and D. Kotturi. Vanishing point detection by line clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(11):1090–1095, 1995.
19. T. Moons, D. Frère, J. Vandekerckhove, and L. Van Gool. Automatic modelling and 3D reconstruction of urban house roofs from high resolution aerial imagery. In *Proc. 5th European Conference on Computer Vision, Freiburg, Germany*, pages 410–425, 1998.
20. S. Noronha and R. Nevatia. Detection and description of buildings from multiple images. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition, Puerto Rico*, pages 588–594, 1997.
21. M. Pollefeys, R. Koch, and L. Van Gool. Self calibration and metric reconstruction in spite of varying and unknown internal camera parameters. In *Proc. 6th International Conference on Computer Vision, Bombay, India*, pages 90–96, 1998.
22. C. Rother. A new approach for vanishing point detection in architectual environments. In *Proc. 11th British Machine Vision Conference, Bristol*, pages 382–391, UK, September 2000.
23. F. Schaffalitzky and A. Zisserman. Planar grouping for automatic detection of vanishing lines and points. *Image and Vision Computing*, 18:647–658, 2000.
24. C. Schmid and A. Zisserman. Automatic line matching across views. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 666–671, 1997.
25. J. A. Shufelt. Performance evaluation and analysis of vanishing point detection techniques. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(3):282–288, March 1999.
26. Hai Tao, Harpreet S. Sawhney, and Rakesh Kumar. A global matching framework for stereo computation. In *Int. Conf. Computer Vision*, pages 532–539. IEEE Computer Society, 2001.
27. C. Taylor, P. Debevec, and J. Malik. Reconstructing polyhedral models of architectural scenes from photographs. In *Proc. 4th European Conference on Computer Vision, Cambridge*. Springer-Verlag, 1996.
28. T. Tuytelaars, L. Van Gool, M. Proesmans, and T. Moons. The cascaded Hough transform as an aid in aerial image interpretation. In *Proc. 6th International Conference on Computer Vision, Bombay, India*, pages 67–72, January 1998.