

TOWARDS FULLY AUTOMATIC GENERATION OF CITY MODELS

Claus BRENNER

Institute for Photogrammetry (ifp), Stuttgart University, Germany
Geschwister-Scholl-Strasse 24D, D-70174 Stuttgart
Claus.Brenner@ifp.uni-stuttgart.de

KEY WORDS: City Models, Building Reconstruction, Feature Extraction, DSM segmentation

ABSTRACT

Once thought of being useful primarily for planning the location of telecommunication antennas, it has become clear in the meantime that three-dimensional city models are of importance in their own right. This paper presents some of our latest results on the reconstruction of building models from laser scanning DSM's and digital ground plans. First, we show how buildings can be reconstructed from ground plans and generalize the standard straight skeleton algorithm. In a second step, we introduce the information which can be obtained from DSM segmentation in order to recover building structures which cannot be inferred directly from the ground plan. The work presented in this paper is actually part of our larger ATOP approach, a new framework for the fully automatic generation of city models.

1 INTRODUCTION

There are numerous applications for three-dimensional city models and nowadays people are becoming increasingly conscious about the technical, environmental and commercial possibilities they offer (see e.g. (Förstner, 1999, Brenner, 1999)). Private companies as well as public authorities realize the vast number of possible applications, including tourism and marketing, architecture and town planning, city climate, noise propagation and environmental research, and new landmark-based navigation systems. It became clear early that photogrammetry is able to provide a means to collect the required three-dimensional information. However, it was evident too that standard photogrammetric systems were not the most efficient solution to measure city models, since they were made to measure single points rather than structured objects. Although research on automatic systems has been going on for over a decade, companies which actually collect three-dimensional models rely almost exclusively on manual data acquisition. Research on building reconstruction can be subdivided into two large classes, namely automatic systems (Henricsson et al., 1996, Henricsson and Baltsavias, 1997, Fischer et al., 1998, Baillard et al., 1999) and semiautomatic systems (Grün and Dan, 1997, Grün and Wang, 1998, Gülch et al., 1999, Brenner, 1999). In summary, we think the following statements reflect the current situation:

- Automatic Systems working solely on the basis of aerial images have shown promising results for some datasets. Often, however, the availability of special data sources has been exploited, such as high-resolution, multiple-overlap or color images. The additional use of a Digital Surface Model (DSM) leads to a more stable detection of buildings, and misdetections due to vegetation can be minimized using a multichannel classification from Color-Infrared (CIR) and DSM data (Haala, 1999). Thus, it has to be noted that the applicability of the methods is constrained to the cases where the mentioned data sources are available, which is seldom the case. Also, most authors admit that their approaches do not work in densely build-up areas, and results for datasets involving several hundred or thousand buildings have not been shown. Therefore, fully automatic systems working on the basis of aerial images are not considered to work reliable enough for practical use, although tremendous progress has been made in this area.
- Automatic Systems working solely on the basis of (Laser-) DSM's have been reported (Brunn and Weidner, 1997). DSM's have the great advantage of representing three-dimensional geometry directly. Also, it is possible to estimate the parameters of planar structures quite accurately when many measured DSM points are available, although the accuracy might be generally lower compared to the measurement from aerial imagery. One of the great disadvantages is the usually low point density of DSM's. This leads to a poor lateral accuracy. Determining the exact location and structure of building walls thus becomes a problem. This might improve when high-resolution DSM's are used (Maas, 1999).
- Combining DSM's with existing ground plans for a fully automatic reconstruction has shown good results, even in urban areas and on datasets with several thousand buildings (Brenner, 1999). Of course, the advantage here is that interpreted information in form of digital ground plans is introduced into the automatic reconstruction process. On the other hand, as buildings become more complicated, the relation between roof and ground plan structure becomes more obscure and can possibly not be exploited anymore. Also, the existence of two-dimensional ground plans must be presumed.

- Semiautomatic systems have been reported (Grün and Wang, 1998, Gülch et al., 1999) which use aerial images and reduce the number of manual interactions required to acquire buildings compared to the fully manual case.
- It can be doubted that automatic systems can achieve success rates comparable to human operators within the next few decades (this statement is due to W. Förstner, (Förstner, 1999)).
- Companies which nowadays collect three-dimensional city models still rely almost exclusively on manual measurement, even though the measured models are not very detailed and/or relatively low accuracies are obtained.

Considering the last two statements, it is our opinion that semiautomatic systems are the only way to serve today's needs. They can be used in large, practical projects, so feedback from users should lead to a fast evolution towards productive systems. They can also be taken as a kernel of a system where fully automatic algorithms can be incrementally tested and integrated. In our opinion, it is not only desirable to have semiautomatic systems which try to reduce the number of operator interactions per building. Rather, one would like to have a system which reconstructs a considerable number of buildings without any operator interference at all. Also, the system should point the operator to those buildings where the automatic reconstruction has possibly failed. Thus, it is our goal to increase the number of buildings amenable to an automatic reconstruction while of course keeping in mind that there will always be a certain percentage left which needs an additional operator-assisted treatment.

In the past, we have reported on our fully automatic building reconstruction system, which relies on a laser scanning DSM and existing ground plans (Brenner and Haala, 1998). The major reconstruction steps of this system were (a) the subdivision of the ground plan into rectangular (2D) primitives, (b) selection of a three-dimensional primitive for each 2D rectangle based on analysis of the DSM (c) estimation of the parameters for each 3D primitive, and (d) assembly of all 3D primitives in order to obtain the model of the entire building. Considering the hierarchy of models in figure 1, the class of buildings that can be reconstructed this way (combined parametric) lies in between the class of parametric and the class of general polyhedral models. Naturally, if one wants to enlarge the number of buildings that can be modelled correctly, a more general building class has to be chosen.




Figure 1: Classes of building models.

2 CONSTRUCTING ROOFS FROM GROUNDPLANS

2.1 Medial axis and straight skeleton

Suppose a ground plan is given in form of a polygon P (a single, simple, planar polygon with no holes). One can imagine that from each polygon segment there emanates one plane with a given slope, inclined towards the interior of P (assume for the moment that all planes have identical slope). The intersections of the planes will yield crest lines (where two planes meet) and points (where three planes meet). When projected down into the plane of P , a planar graph G results. As the graph edges are projections of straight lines in 3D, they are themselves straight lines in 2D.

The *medial axis* $M(P)$ of P is the set of points inside P which have more than one closest point on the boundary of P . At a first glance, it seems that the points covered by a drawing of G and $M(P)$ are identical. However, this is only true for convex P , since the medial axis for polygons with reflex vertices contains curved pieces, which cannot be part of G (Fig. 2(a)). Instead, the graph G can be constructed as *straight skeleton* $S(P)$ (Fig. 2(b)). In order to define the straight skeleton, we can think of a “shrinking process” where the points of the original polygon P move inwards along their angular bisectors until either an “edge event” or a “split event” occurs ((Aichholzer and Aurenhammer, 1995), see also Fig. 2(c,d)). Following these events, one or two new bisectors are generated. Interestingly, the straight skeleton can neither be defined by a distance measure nor can its construction use a divide-and-conquer approach. An $O(N \log N)$ time complexity algorithm (which would be optimal) has not yet been found (Eppstein and Erickson, 1999). However, we shall not be concerned about time complexity here, since in our application, N is typically small.

Given the polygon P and the slopes of all planes, the graph $G = S(P)$ is determined uniquely. The class of roofs defined by the straight skeleton lies in between parametric and polyhedral models. By setting all slopes to zero, prismatic models can be handled. Saddleback roofs are obtained by setting the slope of two roof faces to infinity. However, straight skeleton roofs have only one single eaves height whereas combined parametric roofs may have multiple eaves heights.




Figure 2: (a) Medial axis $M(P)$ for a L-shaped polygon P . (b) Straight skeleton $S(P)$. (c) Edge event (two bisectors intersect). (d) Split event.

2.2 More general roofs

The straight skeleton graph $S(P)$ is a unique solution, but not the only one. In fact, consider the ground plan in figure 3(a) which looks like two separate buildings having a common border. In this case, the straight skeleton reconstruction (Fig. 3(b)) leads to a rather high roof of unusual volume and shape. Intuitively, the reconstruction in figure 3(c) would be considered as being the correct one.




Figure 3: (a) A (ground plan) polygon P . (b) Straight skeleton of (a). (c) Projection of another possible roof.

How can *all* solutions be obtained instead of only the one from $S(P)$? Since the vertices of G are the (projections of the) intersections of three roof planes and edges form the connections between them, the following algorithm can be used: (1) intersect all roof planes to obtain a set of all possible vertices of G , and (2) try all connections between the vertices of G until a valid solution is found. Regarding the first step, the number of intersections of three planes is $\binom{N}{3}$, assuming a ground plan P with N segments. Although that number grows rapidly in the order of $O(N^3)$, for the value range of N considered here this poses no problem. Moreover, since ground plans with large N tend also to have complex, concave shapes, many (often 75% and more) intersection points lie outside P and need not be taken into account. The case is different for the second step. Except for very small N , we cannot hope to solve the search problem in acceptable time by a naive search method such as backtracking. Fortunately, using discrete relaxation, we can reduce search complexity greatly.

2.3 Using discrete relaxation to cut down search space

Consider two intersecting planes e_i and e_j , none of which is vertical, as shown in figure 4. These planes might actually form a concave ($-$) or a convex ($+$) edge when viewed in negative direction of the z axis. We cannot tell from the plane equations which type of intersecting edge is present. However, the edge labeling ($+/-$) directly corresponds to the order of plane labels ($i - j$ or $j - i$).




Figure 4: Intersection of two planes (a) can result in a concave ($-$) or convex ($+$) edge (b). In the projection onto the xy plane (c), the type corresponds to the plane label order.

Now we will turn to the situation at junctions where three planes meet. For line drawing interpretation, there is a classical work of Waltz (Waltz, 1975) who shows that the number of possible labelings for a junction that make sense is much smaller than the number of theoretically possible labelings. Also, using discrete relaxation (Haralick and Shapiro, 1993) it is often possible to obtain a single label for each junction, making a subsequent search unnecessary. In the plane intersection case described above, the situation is slightly different. Since *all* possible intersection points of three planes are generated, the goal is to rule out as many as possible. Any (non-degenerate) intersection of three planes leads to one of the eight cases shown in figure 5. Incoming and outgoing labeled edges of a node constrain the number of possible labelings of the junction.

As an example, figure 6(a) shows the simple case of a rectangular ground polygon P . There are 4 edges of P , and $\binom{4}{3} = 4$ junctions to label. Figure 6(b) shows the situation at junction j_{134} . Since j_{134} is connected to j_{14} , only one




Figure 5: The eight possible labels for junctions where 3 planes meet.

possible label (0) remains. Using the same argument, label 0 remains for junctions j_{234} and j_{124} (Fig. 6(c)). However, in the next iteration step, j_{234} and j_{124} cannot be connected on edge 4 – 2 since their respective arrows do not point towards each other and no valid label remains for those junctions. Thus, in this case only by relaxation without any search, the correct solution is found (j_{134} and j_{123} remain, both with label 0). Of course, in general, several solutions remain after relaxation. A subsequent backtracking search with forward checking can be used to find them. Figure 7 shows results for some cases.




Figure 6: (a) Rectangular ground plan and plane intersection points (junctions). (b) The connection from j_{14} to j_{134} leaves only one valid interpretation. (c) Junctions j_{234} and j_{124} cannot be connected and no valid label remains.

(a)			
(b)	364 (100%)	120 (100%)	220 (100%)
(c)	63 (17%)	50 (42%)	63 (29%)
(d)	29 (8%)	31 (26%)	55 (25%)
(e)	48	65-727	887-67,848
(f)	5602	1647-168,085	269,837-795,654

Figure 7: Examples illustrating search space complexity. (a) Groundplans and all their possible roofs. (b) Total number of junctions. (c) Number of junctions in ground plan P . (d) Junctions remaining after discrete relaxation. (e) Number of search steps when using discrete relaxation (number range when there is more than one solution). (f) Number of search steps without discrete relaxation.

3 HINTS FROM THE DSM

In the previous section, it was assumed that all roof faces have the same inclination. For the algorithms, however, this is not necessary. Changing the inclination of some roof faces will change the topology of the associated graph G . For example, depending on the roof inclinations, the ground plan of figure 6 has two possible associated graphs with junctions $\{j_{134}, j_{123}\}$ and $\{j_{124}, j_{234}\}$, corresponding to two ridge orientations. Would it be, then, possible to simply generate all topologies, fit them to the DSM data (while checking that the fit does not violate the generated topology), and select the one which meets some error criteria best? From the number of possible topologies, we can quickly see that this is not a sensible approach. As just seen, for $N = 4$ we obtain two solutions $\{\{j_{134}, j_{123}\}, \{j_{124}, j_{234}\}\}$. For $N = 5$, there are 5 solutions $\{\{j_{123}, j_{134}, j_{145}\}, \{j_{123}, j_{135}, j_{345}\}, \{j_{124}, j_{145}, j_{234}\}, \{j_{125}, j_{234}, j_{245}\}, \{j_{125}, j_{235}, j_{345}\}\}$ and for $N = 6$, 14 solutions. In general, the number of solutions $s(N)$ satisfies the recursion formula $s(N) = s(N - 1) + s(3) \cdot s(N - 2) + s(4) \cdot s(N - 3) + \dots + s(N - 2) \cdot s(3) + s(N - 1)$ with $s(N) = 0$ for $N < 3$ and $s(3) = 1$. It suffices to see that $s(N) > 2 \cdot s(N - 1)$ and thus $s(N) > 2^{N-3}$, i.e. $s(N)$ grows exponentially.

A more sensible approach is to obtain a first roof hypothesis from a reconstruction assuming identical roof inclinations. In a second step, possible correspondences between the DSM and the reconstructed roof are established. From this, new roof inclinations are derived and the roof reconstruction is repeated. This method has been presented in (Haala and Brenner, 1997). However, these approaches use the DSM only for *measurement* of heights and slopes, whereas the *structure* is taken primarily from the ground plan. On the one hand, this is advantageous when the information content of the DSM is rather low. On the other hand, the disadvantages are:

- roof elements for which there is no hint in the ground plan (such as dormer windows) cannot be reconstructed.
- each segment of P triggers a new roof surface. It is not possible that one roof surface intersects different building borders at different heights.
- ground plans P which are too detailed lead to unnecessary complex roof reconstructions.

In order to improve this, the DSM must also be integrated in the process of finding the roof *structure*, rather than being used for measurement only.

3.1 DSM segmentation

Figure 8 shows the results of several segmentation algorithms for a laser scanner DSM with a ground resolution of one meter. Figure 8(a) shows a segmentation into regions which have normal vectors compatible to the ground plan, and figure 8(b) shows a segmentation based on contours. Finally, figure 8(c) shows a RANSAC-based (Fischler and Bolles, 1981) segmentation into planar regions; in this case, the ground plan is not used except for limiting the segmentation area. In our opinion, approaches which use the ground plan to guide segmentation needlessly narrow the scope for later interpretation steps. Thus, we have chosen the RANSAC-based planar segmentation as input for further processing.




Figure 8: Results of several DSM segmentation algorithms. (a) Normal vector compatibility. (b) Contour based. (c) Segmentation into planar regions.

3.2 A rule-based approach to obtain roof structure

Looking at figure 8(c), one sees that all major parts of the roof result in a segmented region. The task is now to filter the regions and build a final interpretation. Region filtering can be done on the basis of general criteria such as region size, shape and normal vector orientation. Using only constraints of this kind, however, means that no specific building model is part of the reconstruction process. Often, this results in odd roof shapes. Therefore, in our approach we try to combine the segmented regions from figure 8(c) using the ground plan and some rules as follows.

According to their normal vector, regions “along” a ground plan edge can be assigned certain labels (Fig. 9 (left)). For example, if the normal vector of a region is parallel to the normal vector of the ground plan edge, it is called “compatible” (c). Other labels are parallel to previous (p) or next (n) edge, perpendicular to ground plan edge (left (l) and right (r)) and inverse of previous (a) or next (b) edge. Figure 9 (right) shows how some typical building parts would be labeled in the ideal case. Note that one region can get several labels (for example, {n,r,a}). From the sequence of labeled regions along one edge, sequence parts are accepted as follows (< denotes the start, > the end, x^+ one or more occurrences of label x , \star stands for any label):

detected pattern	accepted pattern
$\langle p^+a^+c^+ \rangle$	$\rightarrow \langle p^+a^+ \rangle$
$\langle c^+b^+n^+ \rangle$	$\rightarrow b^+n^+ \rangle$
$\langle p^+ \star n^+ \rangle$	$\rightarrow \langle p^+ \text{ and } n^+ \rangle$
l^+r^+	$\rightarrow l^+r^+$

Additionally, all remaining patterns c^+ are accepted. Regions which have been grouped together during acceptance are inserted into equivalence classes. Each equivalence class undergoes a new plane estimation process. Figure 10(a) shows the regions which have been selected using this algorithm. Note the upper part where a cluster of regions could not be explained by the rules and have been excluded.




Figure 9: Region labeling. Left: Set of possible region labels: compatible (c), previous (p), next (n), left (l), right (r), inverse of previous (a), inverse of next (b). Right: Examples for possible label sequences.

3.3 Putting it together

After accepting planar regions as being part of the roof, one has to actually intersect all regions to obtain a single, connected roof surface. This can be accomplished using a modification of the approach from the previous section. This time, instead of generating a plane for each edge of the ground plan polygon, one plane for each accepted region is generated. Moreover, it is possible to use region adjacency information from the segmentation in order to guide the search process. In figure 10(b), line segments are shown which are computed based on the region plane equations and the region adjacency graph. We cannot use this adjacency information as “hard” constraints for the search, since it is not stable enough. But we can introduce it as a means to reach the correct solution more quickly. Figure 10(c) shows the final roof structure. Note again the upper part, where some of the roof planes are much larger than the segmented regions they are based on. Note also that new roof plane adjacencies not present in the segmentation were generated.




Figure 10: (a) Regions selected by the rule-based approach. (b) Region adjacencies and intersection lines. (c) Final roof structure.

4 RESULTS

Figure 11 shows more results of our approach. In figure 11(a) the situation for a simple saddleback roof is depicted. This roof is handled correctly, albeit a parametric modelling approach or the method described in (Haala and Brenner, 1997) would presumably yield the same result. The situation is different in figure 11(b), where the lower part of the building is actually L-shaped. This results in an additional roof surface for which there is no hint in the ground plan. Also the short horizontal ground plan edge on the left side does not lead to a roof surface. Figure 11(c) shows a similar example, where the simple ground plan suggests a saddleback or hip-roof type building whereas it is L-shaped and has two surfaces which emanate from a ground plan vertex rather than a ground plan edge. Finally, in figure 11(d), the ground plan suggests a fairly complex building (cf. Fig. 7), although from the segmentation, we can see that only two roof surfaces are present. The roof structure is recovered correctly by the algorithm.

5 CONCLUSION AND OUTLOOK

We have presented a new approach for the reconstruction of buildings which relies on ground plans and DSM's from laser scanning. We have shown that this approach is able to reconstruct buildings of fairly complex structure. The entire




Figure 11: Results for several buildings (a–d). (1) Skeleton based on ground plan. (2) Segmented regions. (3) Selected regions from the rule-based approach overlaid with intersection segments. (4) Final reconstructed roof.

process is divided into three stages. In the first stage, a segmentation is used in order to obtain roof surface primitives. In a second step, a rule-based approach decides which segments can be explained by the building model we have chosen. Finally, in the third stage, the roof is built from the primitives that have been accepted, closing any gaps that have been caused by the deletion of unexplainable regions. The algorithm discussed here is only one part of our ATOP (automated




Figure 12: Framework of the ATOP modelling approach.

topology generation for polyhedral objects) modelling approach. In our opinion, past research has too much concentrated on designing single, monolithic systems which are able to reconstruct a wide range of buildings. This may be due to the limited variety of building types in the test datasets which have been widely distributed. In contrast, our goal in ATOP is to provide several “tailored” interpretation modules each of which accepts certain segmentation primitives according to specific rules (Fig. 12). The algorithm presented in this paper as “second stage” is one such interpretation module. Another may be designed exclusively to recover the structure of flat roofs. We can even think of an interpretation module which works according to a rectangle subdivision (Brenner, 1999) or a straight skeleton algorithm. Based on the deviation of the reconstructed surface from the DSM (or image data), the final step of the system consists in selecting the best solution.

REFERENCES

- Aichholzer, O. and Aurenhammer, F., 1995. A novel type of skeleton for polygons. *Journal of Universal Computer Science* 1(12), pp. 752–761.
- Baillard, C., Schmid, C., Zisserman, A. and Fitzgibbon, A., 1999. Automatic line matching and 3D reconstruction of buildings from multiple views. In: IAP, Vol. 32 Part 3-2W5.
- Brenner, C., 1999. Interactive modelling tools for 3D building reconstruction. In: D. Fritsch and R. Spiller (eds), *Photogrammetric Week 99*, Wichmann Verlag, pp. 23–34.
- Brenner, C. and Haala, N., 1998. Rapid acquisition of virtual reality city models from multiple data sources. In: H. Chikatsu and E. Shimizu (eds), *IAPRS Vol. 32 Part 5*, pp. 323–330.
- Brunn, A. and Weidner, U., 1997. Extracting buildings from digital surface models. In: *IAPRS, Vol. 32, Part 3-4W2*, Stuttgart.
- Eppstein, D. and Erickson, J., 1999. Raising roofs, crashing cycles, and playing pool: Applications of a data structure for finding pairwise interactions. *Discrete & Computational Geometry* 22, pp. 569–592.
- Fischer, A., Kolbe, T., Lang, F., Cremers, A., Förstner, W., Plümer, L. and Steinhage, V., 1998. Extracting buildings from aerial images using hierarchical aggregation in 2D and 3D. *Computer Vision and Image Understanding* 72(2), pp. 195–203.
- Fischler, M. A. and Bolles, R. C., 1981. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Comm. of the ACM* 24(6), pp. 381–395.
- Förstner, W., 1999. 3D-city models: Automatic and semiautomatic acquisition methods. In: D. Fritsch and R. Spiller (eds), *Photogrammetric Week 99*, Wichmann Verlag, pp. 291–303.
- Grün, A. and Dan, H., 1997. TOBAGO – a topology builder for the automated generation of building models. In: A. Grün, E. Baltsavias and O. Henricsson (eds), *Automatic Extraction of Man-Made Objects from Aerial and Space Images (II)*, Birkhäuser, Basel, pp. 149–160.
- Grün, A. and Wang, X., 1998. CC-modeler: A topology generator for 3-D city models. In: D. Fritsch, M. English and M. Sester (eds), *IAPRS, Vol. 32 Part 4*, Stuttgart, pp. 188–196.
- Gülch, E., Müller, H. and Läbe, T., 1999. Integration of automatic processes into semi-automatic building extraction. In: *IAPRS, Vol. 32 Part 3-2W5*, München.
- Haala, N., 1999. Combining multiple data sources for urban data acquisition. In: D. Fritsch and R. Spiller (eds), *Photogrammetric Week 99*, Wichmann Verlag, pp. 329–339.
- Haala, N. and Brenner, C., 1997. Interpretation of urban surface models using 2D building information. In: A. Grün, E. Baltsavias and O. Henricsson (eds), *Automatic Extraction of Man-Made Objects from Aerial and Space Images (II)*, Birkhäuser, Basel, pp. 213–222.
- Haralick, R. M. and Shapiro, L. G., 1993. *Computer and Robot Vision*. Addison-Wesley, chapter 17, pp. 19–91.
- Henricsson, O. and Baltsavias, E., 1997. 3-D building reconstruction with ARUBA: A qualitative and quantitative evaluation. In: A. Grün, E. Baltsavias and O. Henricsson (eds), *Automatic Extraction of Man-Made Objects from Aerial and Space Images (II)*, Birkhäuser, Basel, pp. 65–76.
- Henricsson, O., Bignone, F., Willuhn, W., Ade, F., Kübler, O., Baltsavias, E., Mason, S. and Grün, A., 1996. Project AMOBE: Strategies, current status, and future work. In: *IAPRS, Vol. 31, Part B3*.
- Maas, H.-G., 1999. Fast determination of parametric house models from dense airborne laserscanner data. In: *IAPRS Vol. 32, Part 2W1, 5W1, IC5/3W*, Bangkok, Thailand.
- Waltz, D. L., 1975. Understanding line drawings of scenes with shadows. In: P. H. Winston (ed.), *Psychology of Computer Vision*, McGraw-Hill, New York, pp. 19–91.