Towards Fully Automated 3D City Model Generation

C. Brenner, N. Haala & D. Fritsch,
Institute for Photogrammetry, Stuttgart University, Germany

ABSTRACT: Recent developments of efficient software systems have enabled an area covering generation of 3D city models. Still, there is a continuing need for an improvement of these systems to enable a further reduction of human interaction. Nevertheless, the development of new applications including the provision of effective tools for visualization and management of the collected 3D city models now becomes of growing importance. Within the paper our work is exemplary used to give an overview on the current state-of-the-art of 3D city model generation. Considering the reconstruction of building geometry, the paper focuses especially on automatic approaches which integrate 2D ground plan information with digital surface models. A heuristic subdivision of ground plans into rectangles as well as a rule-based reconstruction relying on discrete relaxation are discussed. In the second part further processing tools are discussed based on some application scenarios of the collected data.

1 INTRODUCTION

Three-dimensional city models are usually comprised of a description of the terrain, streets, buildings and vegetation in build-up areas. Building models are an important part thereof, even though it has to be noted that for many applications, additional information is necessary. For example, a faithful representation for virtual reality applications can only be obtained when the texture of the ground, roofs and façades is present and important details like trees, walkways and fences are present (Danahy 1999, Lange 1999).

Interest in three-dimensional building models has raised significantly in the past years. Originally, simulations for the propagation of electromagnetic waves were thought of being one of the major application area. These are used by network operators for the planning of antenna locations. Although there might be an additional need in the near future due to the forthcoming introduction of UMTS networks in Europe, there have been other areas evolving, for example three-dimensional car navigation systems, virtual tourism information systems and visualization for city and building planning or architectural contests.

However, it seems that the most important development-driving factor will not come from those single applications but rather from a principal change which has begun to evolve in the last decade with the widespread use of mobile phones and is characterized by phrases as "ubiquitous computing", "location based services" and "augmented reality". Following this development, location-dependent access to spatial data will become available to any user, (almost) anywhere. Key features will be the accurate localization of the user, the location-dependent processing and transmission of (spatial) data via broadband, wireless networks, availability inside and outside of buildings, and the use of small, yet powerful, mobile devices (Hamburgen et al. 2001).

Therefore, it is an important task to collect spatial data like city models and—in order to keep them useful—to update them in short cycles. It is also a formidable task comparable to the digitization of the road network for navigation purposes. Thus, it becomes clear that any economically viable solution must rely on automatic or at least partially automatic methods.

2 RECONSTRUCTION OF THE BUILDING GEOMETRY

2.1 *Reconstruction Systems*

Photogrammetric methods are well suited for the economic acquisition of 3D city models (Förstner 1999), making it possible to recover the structure as well as the dimensions. On the other hand, classical photogrammetric measurement is mostly point-based, which does not exploit the inherent structure of buildings and thus cannot be optimal economically. Standard automatic image matching techniques, developed originally for the measurement of terrain points, proofed to be not as effective for built-up areas due to the large discontinuities which arise at building borders and roofs. This contradiction has led to substantial research efforts (see e.g. Grün et al. 1995, 1997), and many reconstruction approaches have been proposed. Some coarse classification can be made with regard to the used data sources and the distinction between semiautomatic and fully automatic operation.

Automatic reconstruction from aerial images (e.g. Haala 1994, Henricsson & Baltsavias 1997, Fischer et al. 1998, Baillard & Zisserman 1999) has shown promising results, however one has to note that often special image material has been used which is not available in general, for example large scale, multiple overlap, or colour images, or additional height models. Even then, the reliable extraction of buildings in densely build-up areas has not been demonstrated yet.

Automatic systems working solely on the basis of digital surface models (DSMs) acquired by laser scanning have been reported (Brunn & Weidner 1997, Maas 1999). Since DSMs represent the geometry of the surface directly, they have advantages with regard to automated interpretation. They used to suffer from the low density of measured points, however this is no technical restriction anymore, since a density of several measured points per square meter is attainable. Also, the recording of multiple return pulses and intensity values is feasible nowadays, opening up new possibilities for automated segmentation techniques.

Semiautomatic approaches have been reported both for image- and DSM-based systems. They can be divided into approaches which model buildings from a fixed set of volumetric primitives which are combined (e.g. Gülch et al. 1999, Brenner 1999) and approaches which build the topology of the surface directly (e.g. Grün & Wang 1998).

2.2 *Reconstruction using ground plan segmentation*

A high percentage of buildings can be modeled using a small number of building primitives like flat boxes, boxes with saddleback and hipped roofs and other geometric primitives like cylinders and cones. Even if one considers only primitives based on rectangular ground shapes, still the majority of buildings can be modeled.

Thus, one approach to reconstruct buildings when 2D ground plans are available can be sketched as follows:

- Try to infer (in 2D) from the ground plan how the building can be subdivided into primitives.
- Select each primitive based on additional information, e.g. an aerial image or a DSM. Determine the dimensions by a measurement process.
- Assemble all primitives to obtain a single body representing the building.

For example, for a simple L-shaped building the first step would select two rectangles to cover the 2D ground plan, the second step would select hipped roofs on each of those rectangles and estimate the eaves and ridge heights and the third step would finally merge both volumetric primitives.

This approach has been used by the authors in several projects, combining 2D ground plans with DSMs from laser scanning (Brenner 1999). Its main advantage is that by the use of digitized ground plans, interpreted information is "injected" into the reconstruction process which makes it relatively simple and reliable. On the other hand, since the selection of primitives is guided by analyzing the ground plan, the final reconstruction is strongly coupled to the ground plan shape. For example, roof structures like dormer windows will not appear in the result as long as there is no corresponding hint in the ground plan. This situation can be improved by interactive modeling tools, however.

A second disadvantage is, that while most simple buildings lend themselves well to a subdivision into primitives, there are always buildings present which cannot be modeled properly or for which a subdivision into primitives is not very natural and thus modeling becomes involved. In those cases, other approaches which allow to specify the roof topology directly would be more desirable.

2.3 *Reconstruction using tree search and DSM segmentation*

2.3.1 *Ground plans and simple roofs*

We consider a ground plan given as a closed polygon P consisting of n linear segments p_i and vertices v_{12}, v_{23}, ..., $v_{n-1,n}$, $v_{n,1}$ (where P does not contain any inner polygon), see Figure 1. Based on P, a simple roof can be constructed which consists of planar roof faces Π_i which intersect with the building walls in the eaves P_i and eaves points $V_{i,i+1}$. The intersections of the roof faces yields ridges (E_{12}, E_{24}) and grooves. Projected down into the plane the corresponding points and lines (indicated by lowercase letters) are obtained. If the original ground plan is omitted, the remaining points and segments form a *planar graph G*. It consists of leaves v_{ij}, inner nodes v_{ijk} and edges e_{ij}. If the indices are viewed as sets, then the edge e_A connects nodes v_B and v_C iff $A=B\cap C$. The graph G is connected and without cycles which means it is a tree.

Simple roofs such as the one in Figure 1 can be recovered from the corresponding planar Graph G by adding a ground and eaves height. The graph G itself can be obtained from the ground plan P by computing the *straight skeleton S(P)*. For convex ground plans, the straight skeleton *S(P)* is identical to the *medial axis transform M(P)*. It is different for ground plans containing concave nodes, in which case *M(P)* contains parabolic segments. The construction of the straight skeleton is a well-known process in geometry but has received attention recently (Aichholzer et al. 1995). Interestingly, no time-optimal algorithm has been found so far (Eppstein & Erickson 1999).

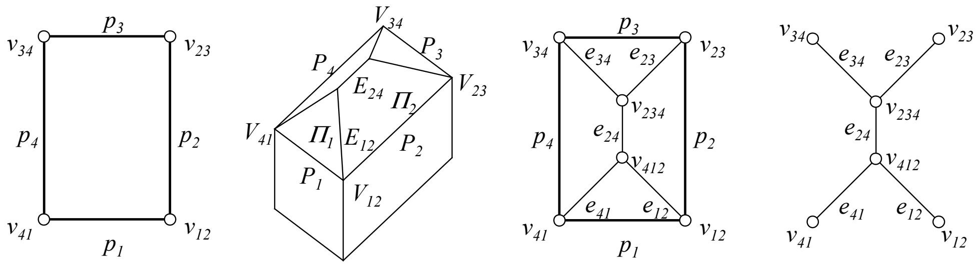

Figure 1. From left to right: ground plan, corresponding simple roof, projection into the plane and corresponding graph (which is the skeleton *S(P)* of *P*).

2.3.2 *From construction to search*

The straight skeleton *S(P)* can be constructed from a given ground plan P unambiguously. However, it is often not the only roof which can be raised on P.

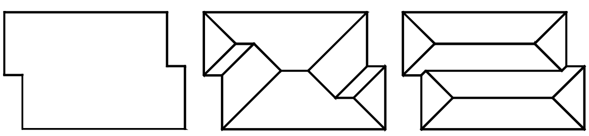

Figure 2. From left to right: ground plan P, P together with graph G corresponding to the straight skeleton and another possible graph G'.

Figure 2 shows an example. In order to obtain all possible solutions, the constructive algorithm has to be replaced by a search procedure. If a roof face Π_i is defined emerging from each eaves edge P_i, extending inwards, then all possible vertices V_{ijk} can be constructed by intersecting the corresponding planes Π_i, Π_j and Π_k. The graph G is then obtained by searching all possible edge combinations between those vertices.

Considering the first step, the intersection of every three roof planes has a time complexity of n over 3, or O(n^3), i.e. polynomial in the number of ground polygon segments. However, the subsequent search has exponential complexity and is usually not solvable except for very small values of n.

Fortunately, one can easily do better by labeling edges and nodes. This is similar to the approaches for the interpretation of line drawings – a task which has been a topic of artificial intelligence research from its very beginning (Waltz 1975). Indeed, edges e_{ij} are the projection of the intersection of two planes where (depending on the selection of half-planes) either a concave (–) or a convex (+) intersection is present (Fig. 3). Thus, any edge e_{ij} allows two different interpretations. Similarly, it turns out that each node v_{ijk} allows for eight different interpretations, where each interpretation defines the type (+/–) as well as the direction of all incident edges.

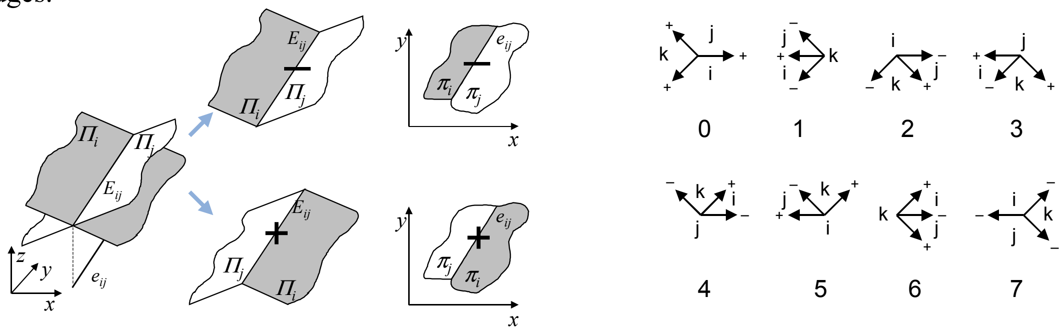

Figure 3. Possible interpretations of edges (left) and nodes (right).

Thus, a constrained tree search can be used to cut down search space. A still more dramatic reduction can be obtained when a discrete relaxation is used beforehand in order to reduce the number of possible edge and node interpretations. Figure 4 shows some examples. In general, discrete relaxation propagates constraints on the number of possible interpretations starting from the ground polygon nodes towards inner nodes (since the interpretation of ground polygon nodes $v_{i,i+1}$ and edges $e_{i,i+1}$ is determined by the shape of the ground plan).

(a)

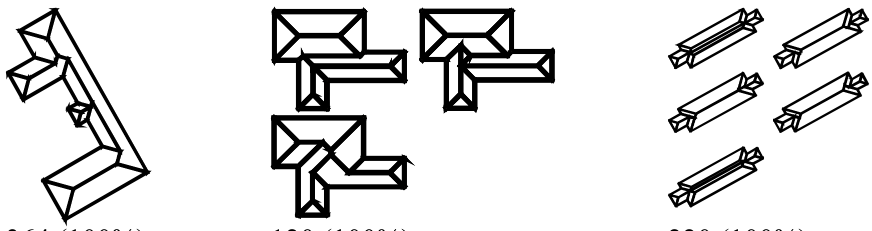

(b)	364 (100%)	120 (100%)	220 (100%)
(c)	63 (17%)	50 (42%)	63 (29%)
(d)	29 (8%)	31 (26%)	55 (25%)
(e)	48	65 – 727	887 – 67848
(f)	5602	1647 – 168085	269837 – 795654

Figure 4. (a) Ground plans and possible roofs found by the search. (b) Maximum number of nodes (which is n over 3). (c) Number of nodes which are inside the ground plan. (d) Remaining nodes after discrete relaxation. (e) Required search steps when discrete relaxation is used (range when more than one solution exists). (f) Required search steps without discrete relaxation.

2.3.3 *Incorporating digital surface models*

So far, the generation of roofs based solely on ground plans has been discussed. However, in order to obtain faithful reconstructions of real world objects, not only the dimensions but also the structure has to be inferred from measurements – in our case DSMs from laser scanning.

There are numerous techniques to segment DSMs into meaningful regions, for example region growing, line grouping (Jiang & Bunke 1994) or robust estimation techniques like RANSAC (Fischler & Bolles 1981). Usually, a DSM segmentation into *planar* regions is sensible for roof reconstruction purposes, although with higher DSM point densities, one might also try to recover higher order surfaces. Figure 5 shows an aerial image of a complex building and a planar region segmentation obtained from a 1 m grid DSM. As can be seen, all important regions are captured, however there are many other small regions present. In order to reconstruct the roof's surface, those have to be removed or merged with other regions. One could use general criteria for this such as the size or shape of the regions. However, better results can be obtained when a building specific model is used as in the following rule-based approach.

Figure 5. Left: aerial image of a building with complex roof. Right: regions obtained from a planar DSM segmentation.

2.3.4 *A rule-based approach for the acceptance of regions*

If we compare the normal vector of a segmented region with the orientation of an adjoining ground plan edge, we can derive a certain relationship. For example, for the standard case where a roof "evolves" from the corresponding ground plan edge (the normal vector is perpendicular to the edge, pointing outwards), we would label it as "compatible" (c). Similarly, there are labels for compatibility with the previous (p) and next (n) edge, with opposite vectors of the previous (a) or next (b) edge or perpendicular to the current edge to the left (l) or right (r). Each region can have several labels, where with rectangular ground plans, {l,p,b} and {r,n,a} are typical label sets (see Fig. 6, left). For example, it can be seen that dormer windows result in the typical sequence c, l, r, c. Regions are accepted based on the rules shown in Figure 6, right, where "<" stands for the start and ">" for the end of a pattern along one edge, „x^+“ stands for one or more occurrences of the pattern "x" and "*" stands for any pattern. The order in which the rules are applied is of importance. All remaining regions along an edge labeled with "c" are accepted. Figure 7 shows the result of the rule application on the segmentation from Figure 5.

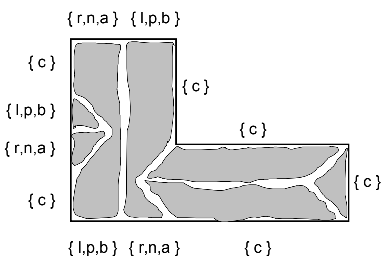

Detected pattern		Accepted Pattern
$<p^+a^+c^+$	$\rightarrow$	$<p^+a^+$
$c^+b^+n^+>$	$\rightarrow$	$b^+n^+>$
$<p^+ * n^+>$	$\rightarrow$	$<p^+$ und $n^+>$
l^+r^+	$\rightarrow$	l^+r^+

Figure 6. Left: an example of segmented regions, ground plan and resulting labeling. Right: rules used to accept a subset of regions.

In order to derive a closed roof topology, the search method of section 2.3.2 can be modified. Although it has been presented in conjunction with a ground-plan based reconstruction, it only uses the intersection of planes. Instead of implicit planes emerging from their corresponding ground plan edges, the planes defined by accepted regions can be used. The result of this step is shown in Figure 7, right. It can be seen that the gap in the right area of the building (Fig. 7, left) which is due to non-accepted regions has been closed. While this reconstruction is not fully correct, it is in accordance with the model as given by the rules. In contrast to the reconstruction of Figure 4, the rooftop in the middle part of the building as well as the dormer window in the left part are captured.

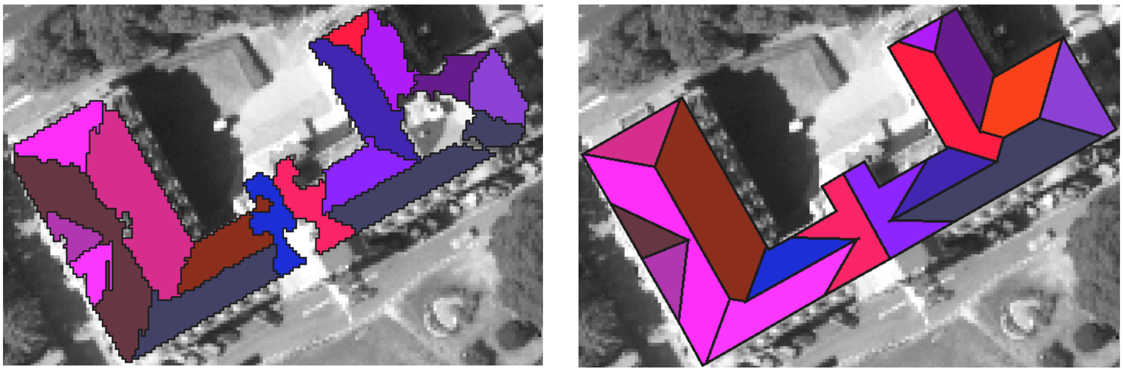

Figure 7. Accepted regions (left) and final reconstructed topology (right).

3 APPLICATION OF TERRESTRIAL IMAGERY

Similar to the methods discussed above, almost all current systems apply airborne data for the collection of 3D city models. Of course data capture is also feasible based on terrestrial images. Commercially available software tools allow for 3D measurement at high accuracies, nevertheless close range techniques for architectural photogrammetry currently are too time consuming for an area covering data collection. Airborne data is more or less equivalent to terrestrial images if geometric data capture is aspired, but the integration of terrestrial imagery is mandatory for applications like texture mapping. Thus, the lack of tools for efficient processing of this type of data defines a bottleneck in the current data flow.

Figure 8: 3D City Model of Heidelberg

3.1 *Texture mapping*

The appearance of powerful tools for interactive visualization of synthetic landscapes in computer vision was one of the driving forces motivating the great efforts for 3D city model collection. Since one of the most important applications of virtual city models is the generation of realistic visualizations, a proper representation of geometry and texture has to be provided for each building. If aerial images are used for geometric data capture, the texture for each building surface is already available as a by-product. In each case, the exterior orientation of the imagery is either already available or can be determined easily by standard software. Since this information permits a simple back-projection of the reconstructed 3D building to the aerial image, the corresponding image patch for each visible part of a building is available after this step and can be used for texture mapping.

An example of a visualization solely based on aerial image texture is given in Figure 8. The geometry of this virtual city model was collected by the approach discussed in section 2.2.. Presentations similar to this example are sufficient if nadir or oblique views are generated at medium distances in order to depict a smaller district or at least a couple of buildings. If the observer approaches even closer during the generation of a virtual walk-through, texture mapping exclusively based on airborne images is not adequate any more. Especially the geometric resolution of image patches depicting vertical walls is insufficient due to the disadvantageous viewing angle in airborne data. Since these façades are one of the most important elements to be visualized, the additional integration of terrestrial images is inevitable to overcome this problem in such application scenarios.

Except for industrial applications automatic procedures for tie and control point measurement from terrestrial images are not available. This is the main reason why the reconstruction of the exterior orientation for this data is more time consuming compared to the processing of aerial imagery. Thus, in the current working flow this effort is reduced by a simplified approach, which applies a manual mapping of the façades. A simple software tool allows for the manual selection of a single building face from the terrestrial image, afterwards the resulting image patch is rectified and attached to the corresponding 3D polygon of the building. The visible façades of buildings as depicted in Figure 9 were generated by this procedure.

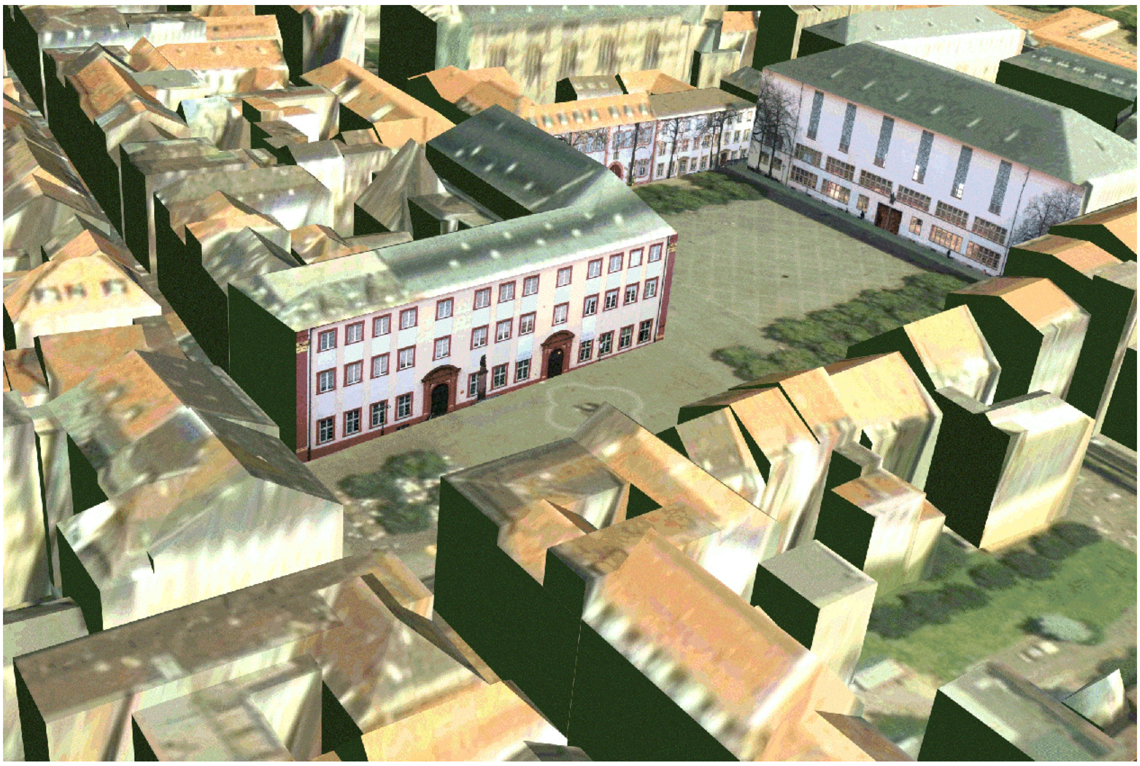

Figure 9: Visualization based on combination of aerial images and terrestrial texture

In addition to the considerable amount of human interaction, the main disadvantage of this approach is the need to repeat the manual texture mapping after changes of the reconstructed building geometry. If a building polygon is changed, eliminated or split, its former link to the rectified image patch is lost and thus has to be restored by the human operator again. These changes of building geometry are for example necessary if the reconstruction is refined by additional measurements. In most cases these changes will be required to enable real-time visualization. For these applications the amount of data to be displayed usually has to be reduced based on multi-scale representations, thus a number of different geometric representations of the same object has to be generated and textured. At present multi-scale techniques are mainly used for terrain rendering (Heckbert & Garland 1997), but current research effort also aims on modifications of these approaches for the visualization of buildings (Decoret et al. 1999).

Thus, in the medium term, automatic tools for an efficient georeferencing of terrestrial images must be available to enable automatic texture mapping. Even more important, a major demand for orientation techniques will also emerge from applications of virtual city models in the context of location based services. One approach to realize these services is the presentation of virtual city models by augmented reality techniques.

3.2 *Augmented reality*

One of the most promising applications of virtual city models is their integration into personal navigation systems. In this case a 3D visualization of the environment can replace 2D map-like representations in order to enable a more realistic navigation. Even more important, supplementary information can be presented very intuitively by augmented reality techniques. Context dependent information is fitted to the real objects being viewed and presented to the user by devices like data-glasses or head-mounted displays. Augmented reality is for example used in technical applications to replace a handbook. Here what a technician actually sees is complemented by an overlay of the relevant fitting instructions for the component he is actually looking at. Similar scenarios are also feasible for city information systems where the visible environment is enriched by information relevant for each building.

The virtual city model presented in Figures 8 and 9 was generated within the Deep Map project, which aimed on the development of a mobile tourist information system for the city of Heidelberg. Within this project a mobile tourist-information-system was developed to enable virtual walks, but also the realization of on site queries on thematic information like opening hours of museums or the generation and overlay of historic views is aspired (Coors et al. 2000). A similar system, helping a user to navigate through a build-up area is also described by Höllerer et al. (1999). Using a head mounted display, the names of buildings are provided to the user based on his actual field of view. By pointing to the buildings additional information is made accessible by an integrated wireless access to the internet. This so-called telepointing feature is also realized within the project NEXUS (Fritsch et al. 2000). In this environment the head-mounted display is replaced by an image of the users environment. This image can for example be captured by a camera integrated into a small hand-held display. If the exterior orientation of the image is available, queries on visible objects can be realized by pointing to regions of interest directly on the display.

3.3 *Orientation of terrestrial images for location based applications*

One option to determine the required position and orientation information for the collected imagery is direct georeferencing. A commercial system, which integrates geocoded image sequences together with further information into electronic city maps is described by Sood & Fahrenhorst (1999). Nevertheless, the quality of the collected exterior orientation based on GPS measurement for their configuration is not sufficient for telepointing or texture mapping. In contrast to this Bosse et al. (2000) describe a system for the collection of georeferenced terrestrial images at high accuracies in urban areas using DGPS/INS measurement. In their application these images are used for the subsequent collection of building geometry, thus the required accuracies on the system hardware are relatively high. Alternatively, the demanded accuracy of the measured camera position and orientation can be reduced if a 3D model of the buildings at the site is already available. In that case the terrestrial imagery can be aligned to the reconstructed buildings e.g. as described by Jaynes (1999).

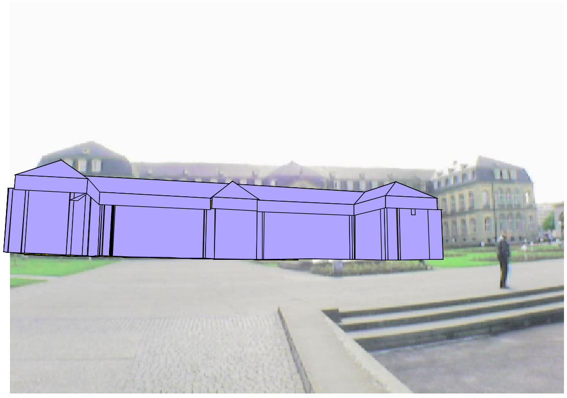

Figure 10: Projection of building based on GPS and digital compass measurement

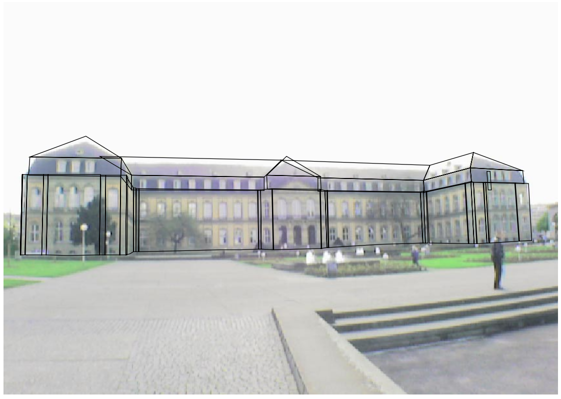

Figure 11: Projection of building based on refined orientation.

In our current system a digital camera is combined with a DGPS receiver, a digital compass and a tilt sensor module. The positional accuracy of 1 – 3 m as well as the orientation accuracy of 1° – 2° of these low-cost components already enables a selection of visible buildings and a initial transformation of the building geometry to the collected terrestrial image (see Fig. 10). Telepointing, as it is required within our NEXUS project is feasible in most cases, whereas texture mapping requires a refinement of the camera pose based on spatial resection (Fig. 11). At the moment, this information is provided by manual measurement, however current work aims on the automation of this process.

4 CONCLUSIONS

Within this paper two different approaches have been discussed for the geometric building reconstruction. The first one uses volumetric primitives. The second builds the roof topology based on DSM segmentation, a set of rules and a search procedure.

In general, nowadays semi-automatic tools are available, which have already proven their capability for an area covering capture of building geometry and DTM data in larger projects, significant improvement has also been reported for automatic systems. Therefore, in addition to the further improvement and automation of algorithms for the capture of virtual city models, the development and promotion of new applications becomes of growing importance. As described within the second part of this paper, these applications presume the (semi-)automatic orientation of terrestrial imagery in an urban environment, thus the development of appropriate tools will be one of our research goals in the near future.

REFERENCES

Aichholzer, O., Aurenhammer, F., Alberts, D., Gärtner, B. (1995). A Novel Type of Skeleton for Polygons, Journal of Universal Computer Science 1(12), 752 – 761.

Baillard, C., Zisserman, A. (1999). Automatic reconstruction of piecewise planar models from multiple views. Proc. IEEE Conference on Computer Vision and Pattern Recognition, 559-565.

Bosse, M., De Couto, D., and Teller, S. (2000). Eyes of Argus: Georeferenced Imagery in Urban Environments. GPS World April, 20-30.

Brenner, C. (1999). Interactive modelling tools for 3D building reconstruction. In: D. Fritsch & R. Spiller (Eds.), *Photogrammetric Week '99*, p. 23–34, Wichmann Verlag.

Brenner, C. (2000). Dreidimensionale Gebäuderekonstruktion aus digitalen Oberflächenmodellen und Grundrissen, Ph.D. Thesis, Deutsche Geodätische Kommission, Vol. C 530 (in German).

Brunn, A., Weidner, U. (1997). Extracting Buildings from Digital Surface Models. In: IAPRS, Vol. 32, Part 3-4W2, Stuttgart.

Coors, V., Huch, T., and Kretschmer, U. (2000). TouriSTAR-ein mobiles AR System in einer urbanen Umgebung. Visualisierung raumbezogener Daten: Methoden und Anwendungen, Bd. II. 191-207. Münster, Institut für Informatik, Uni Münster.

Danahy, J. (1999). Visualization Data Needs in Urban Environmental Planning and Design. In: D. Fritsch & R. Spiller (Eds.), *Photogrammetric Week '99*, 351–365, Wichmann Verlag.

Decoret, X. S. G., Silion, F., and Dorsey, J. (1999). Multi-layered impostors for accelerated rendering. EUROGRAPHICS '99 18[3].

Eppstein, D., Erickson, J. (1999). Raising Roofs, Crashing Cycles, and Playing Pool: Applications of a Data Structure for Finding Pairwise Interactions, Discrete & Computational Geometry 22, 569 – 592.

Fischer, A., Kolbe, T., Lang, F., Cremers, A., Förstner, W., Plümer, L., Steinhage, V. (1998). Extracting Buildings from Aerial Images using Hierarchical Aggregation in 2D and 3D. Computer Vision and Image Understanding 72(2), 195-203.

Fischler, M. A., 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), 381 – 395.

Förstner, W. (1999). 3D-City Models: Automatic and Semiautomatic Acquisition Methods. In: D. Fritsch & R. Spiller (Eds.), *Photogrammetric Week '99,* 291–303, Wichmann Verlag.

Fritsch, D., Klinec, D., and Volz, S. (2000). NEXUS - Positioning and Data Management Concepts for Location Aware Applications. Proc. of the 2nd International Symposium on Telegeoprocessing. 171-184.

Grün, A., Baltsavias, E., Henricsson, O. (Eds.) (1997). *Automatic Extraction of Man-Made Objects from Aerial and Space Images (II)*, Birkhäuser, Basel.

Grün, A., Kübler, O., Agouris, P. (Eds.) (1995). *Automatic Extraction of Man-Made Objects from Aerial and Space Images*, Birkhäuser, Basel.

Grün, A., Wang, X. (1998). CC-modeler: A topology generator for 3-D city models. In: D. Fritsch, M. Englich, M. Sester (Eds.), IAPRS, Vol. 32, Part 4, Stuttgart, 188-196.

Gülch, E., Müller, H., Läbe, T. (1999). Integration of automatic processes into semi-automatic building extraction. In: IAPRS, Vol. 32, Part 3-2W5, Munich.

Haala, N. (1994). Detection of Buildings by Fusion of Range and Image data, IAPRS Vol. 31 Part 3, Munich.

Hamburgen, W. R., Wallach, D. A., Viredaz, M. A., Brakmo, L. S., Waldspurger, C. A., Bartlett, J. F., Mann, T., Farkas, K. I. (2001). Itsy: Stretching the Bounds of Mobile Computing, IEEE Computer, 18(4), April, p. 28–36.

Heckbert, P. S. and Garland, M. (1997). Survey of Polygonal Surface Simplification Algorithms. SIGGRAPH '97 Course Notes Multiresolution Surface Modeling. 8-38. Los Angeles.

Henricsson, O., Baltsavias, E. (1997). 3-D Building Reconstruction with ARUBA: A Qualitative and Quantitative Evaluation. In: Grün, A., Baltsavias, E., Henricsson, O. (Eds.), *Automatic Extraction of Man-Made Objects from Aerial and Space Images (II)*, 65-76.

Höllerer, T., Feiner, S., Terauchi, T., Rashid, G., and Hallaway, D. (1999). Exploring MARS: developing indoor and outdoor user interfaces to a mobile augmented reality system. Computers & Graphics 23, 779-785.

Jaynes, C. (1999). View Alignment of Aerial and Terrestrial Imagery in Urban Environments. ISD'99. 3-19. Springer-Verlag.

Jiang, X. Y., Bunke, H. (1994). Fast Segmentation of Range Images into Planar Regions by Scan Line Grouping, Machine Vision and Applications 7(2), 115 – 122.

Lange, E. (1999). The Degree of Realism of GIS-Based Virtual Landscapes: Implications for Spatial Planning. In: D. Fritsch & R. Spiller (Eds.), *Photogrammetric Week '99*, 367–374, Wichmann Verlag.

Maas, H.-G. (1999). Two Algorithms for Extracting Building Models from Raw Laser Altimetry Data. ISPRS Journal of Photogrammetry and Remote Sensing 54, 153-163.

Sood, R.-A. and Fahrenhorst, C. (1999). Geocoded Image Sequences. Photogrammetric Week 99. 313-316. Wichmann Verlag.

Waltz, D. L. (1975). Understanding line drawings of scenes with shadows. In: P. H. Winston (Ed.), Psychology of Computer Vision, McGraw-Hill, New York, 19 – 91.