

FAÇADE MODELLING FOR HISTORICAL ARCHITECTURE

Jan Böhm, Norbert Haala, Susanne Becker

Institute for Photogrammetry, Universitaet Stuttgart
Geschwister-Scholl-Str. 24D, D-70174 Stuttgart, Germany
Forename.Lastname@email.address

KEY WORDS: Three-dimensional, Point Cloud, Urban, LIDAR, Modelling

ABSTRACT:

The paper presents an approach to efficiently increase the geometric detail of complex building facades by terrestrial laser scanning (TLS). Such applications are frequently required for documentation and presentation of architectural heritage. The modeling of the building geometry is based on a coarse to fine approach. Thus, an approximate 3D building model like it is typically available from airborne data collection provides a-priori information to support the subsequent processing of the terrestrial data. Since the generation of a true CAD description of facade primitives such as windows, doors and balconies is supported by a close integration of such existing building models, the data collection from the scratch can be avoided and the modelling process is facilitated. Thus, after the alignment of the collected LIDAR point clouds to the given 3D model, window structures are extracted and integrated to a refined CAD model. These structures can be further refined by edges extracted from additional façade images. Alternatively to such an explicit 3D representation of façade structures, which is realized by an approach based on cell decomposition, a planar mapping of the LIDAR point clouds to generate a 2.5D representation can be used. By these means a refined visualization of the façade geometry is feasible based on standard computer graphic techniques like bump or displacement mapping.

1. INTRODUCTION

A large portion of cultural heritage applications focuses on architectural heritage. There, besides the coarse overall geometry of a building, the facades are of special interest. Fine geometric details of windows, balconies, stonework and ornaments give each building its individual character. It is these features that make a building worthy for heritage recording. Since façade details are three-dimensional in nature, three-dimensional data acquisition is required for accurate recording. Modern laser scanners are able to acquire densely sampled point clouds of a façade consisting of several million points. Using state-of-the-art meshing techniques, though not without problems, these point clouds can be converted to polygonal meshes, immediately suited for graphic rendering. While dense polygon meshes might be a suitable representation for applications in computer graphics, in the world of geomatics this representation lacks desirable features. Let alone the size of the data makes it unrealistic to derive such a representation for a larger number of buildings. Furthermore a dense polygon mesh is not grouped into meaningful entities, such as different architectural structures, for example wall, roof, and so on.

Deriving such a grouping leads to the task of modelling, where geometric primitives are fit to segmented portions of the data. For TLS data this is typically an interactive process, involving the manual segmentation of the point cloud and the selection of appropriate primitives. This approach is currently applied for the representation of industrial scenes, where objects can be described with a reduced set of geometric primitives and graphical realism is of little interest. For the representation of general architectural scenes this approach of modelling does not seem suitable. For one the work load of manual interaction limits the approach to single buildings or even portions of a facade. Secondly most modelling environments for TLS data are limited to simple surfaces types such as ruled surfaces,

seldom higher order curved surfaces are used. This set of geometric primitives is not suitable for the representation of the highly detailed facades of historical buildings.

In this paper we explore two different ways to model the façade of a historical building. Both approaches rely on the existence of a coarse building model that describes the overall geometry of the building in a polyhedral approximation. This coarse representation can for example be already available from airborne data collection. The first approach attempts to derive a true CAD description of facade primitives such as windows, doors and balconies. The approach is based on a cell decomposition process. The initial portion of the point cloud covering the part of the façade which is to be modelled is segmented into cells belonging to the main plane of the facade and out-of plane cells which belong to windows or other structures on the façade. Geometric primitives are then fitted to the out-of-plane cells. By these means an effective segmentation of the façade into meaningful entities can be achieved. Geometrically it is only limited by the number of primitives implemented.

The second approach is to convert the point cloud generated by the scanner into a LASERMAP. The term is composed from two terms describing the source of the data, a laser scanner, and the use of the data, as a displacement or bump map. The LASERMAP is a re-interpolation of the point cloud into a regular raster centred on the planar polygon describing a single portion of the façade very similar to a Digital Elevation Model. Each computed LASERMAP is then mapped onto the faced polygon for rendering. It can either be directly used as a displacement map, when a suitable rendering engine is available or a bump map can be computed using simple normal vector computation from a height field. This approach preserves nearly all details of the original point cloud and can represent all two-and-a-half-dimensional features of a façade. It is ideally suited

for visualization, direct interaction and virtual reality applications. Alternatively, combinations of the two approaches are possible, where CAD primitives are fit to the data where suitable primitives are available and certain geometric constraints have to be met, while LASERMAPs are used otherwise.

2. GENERATION OF CAD DESCRIPTION FOR REFINED 3D BUILDING MODELS

In order to combine the terrestrial laser data with a pre-existing building model, the different data sets have to be registered first. Several methods are available to provide this registration, ranging from direct georeferencing, to manual and automatic alignment.




Figure 1: 3D point cloud from laser scanning aligned with a virtual city model.

Figure 1 shows an overlay of the measured TLS data and a coarse CAD model based on the registration and georeferencing approach for terrestrial laser data described by Schuhmacher and Böhm (2005). The point cloud was acquired by a Liеча HDS 3000, while the 3D building representations were extracted from an existing area covering 3D city model. This data set is maintained by the City Surveying Office of Stuttgart and gives a quality and amount of detail, which is typical for 3D models from airborne data collection.

In the following, the geometric refinement of the building façade is realised by a two-step approach. First, the LIDAR point clouds are used to decompose the given building model into 3D cells to additionally represent façade structures like windows and doors. This cell decomposition can be used very effectively to represent building models at multiple scales (Haala et al 2006). Thus, it is refined in a second step by photogrammetric analysis of the images. By these means, the amount of detail is further increased for the window frames while profiting from the higher resolution of the image data.

2.1 Point cloud segmentation

For model refinement the available coarse 3D building with its flat front faces is decomposed into suitable façade cells. For this purpose, planar delimiters have to be derived by extracting window borders from the measured LIDAR points. Thus, as a first step of the LIDAR based refinement of the building façade, suitable 3D point measurements are selected by a simple buffer operation. While assuming that the façade can be described sufficiently by a relief, the vertical distances between the measured 3D laser points and the given façade polygon can be used to generate a 2½ D representation, or can even be

interpolated to a regular grid. Thus, further processing like the following segmentation is simplified considerably by such a mapping of the 3D points against the given façade plane.




Figure 2 Detected edge points for a planar façade polygon at horizontal and vertical window structures.

Usually, fewer 3D points are measured on the façade at window areas. This is due to specular reflections of the LIDAR pulses on the glass or points that refer to the inner part of the building. For this reason, our point cloud segmentation algorithm detects window edges by searching for areas with no point measurements. An exemplary result is given in Figure 2. In principle, areas with no point measurements can also result from occlusions. However, this is avoided by using point clouds from different viewpoints. In that case, occluding objects only reduce the number of LIDAR points since a number of measurements are still available from the other viewpoints. Our segmentation process differentiates four types of window borders: horizontal structures at the top and the bottom of the window, and two vertical structures that define the left and the right side. As an example, the edge points of a left window border are detected if no neighbour measurements to their right side can be found in a pre-defined search radius, depending on the expected point density at the façade plane.

2.2 Façade Structure from TLS

The actual representation of the façade structure is based on cell decomposition. Similar to constructive solid geometry (CSG) this type of representation describes complex solids by a combination of relatively simple, basic objects in a bottom up fashion. In contrast to CSG, which combines simple primitives by means of regularized Boolean set operators, decomposition models are limited to adjoining primitives. Since the basic primitives must not intersect, they are thus ‘glued’ together to get the final model. In this sense, cell decomposition is similar to a spatial occupancy enumeration, where the object space is subdivided by non overlapping cubes of uniform size and orientation. However, cell decompositions are based on a variety of basic cells, which may be any objects that are topologically equivalent to a sphere i.e. do not contain holes. This allows for a simplified combination of the respective elements, while the disadvantages of exhaustive enumeration like large memory consumption and the restricted accuracy of the object representation can be avoided.

As already mentioned, the flat front face of a 3D building is separated into 3D cells. Each 3D cell represents either a homogeneous part of the façade or a window area. In order to generate such window cells, the extracted edge points are used to determine window borders by estimating horizontal and vertical lines from non-isolated edge points. These boundary

lines are depicted together with the used edge points in Figure 3.




Figure 3. Detected horizontal and vertical window lines.

In the following, they are used to decompose the building façade in suitable cells. Each of these cells either defines a homogeneous part of the façade or a window area. The discrimination between window and façade cells can be realised based on the availability of measured LIDAR points. After this classification window cells, i.e. areas with no point measurements, can be eliminated from the façade. Within this step, neighbourhood relationships and constraints concerning the simplicity of the resulting window objects can be realized. As an example, uncertain cells are classified depending on their neighbours in order to align and adapt proximate windows in horizontal and vertical direction and to guarantee convex window objects.




Figure 4. Refined facade of the given building model.

As it is depicted in Figure 4, the façade geometry is modelled by eliminating the classified window cells from the existing coarse building model. For this purpose, a plane parallel to the facade at window depth is determined from LIDAR points measured at the window crossbars. The window cells are then carved out from the building model at this window depth.

While the windows are represented by polyhedral cells, also curved primitives can be integrated in the reconstruction process as demonstrated by the round-headed door of the building. Furthermore, our approach is not limited to the modelling of indentations like windows or doors and can be used to add details as protrusions to the façade.

2.3 Refined Façade Structure from Stereo Image Analysis

Usually, the limited point sampling distance of TLS restricts the amount of geometric detail which can be provided by the reconstruction process. Thus, additional 3D information is optionally captured from stereo images. As an example window crossbars can be reconstructed from linearly shaped façade detail. These objects are provided from edge points, which are extracted and matched by a suitable process. First, a standard Sobel filter is used for edge extraction, which are thinned and split into straight segments. Afterwards, the resulting 2D edges of the stereo image pair are matched to generate the required 3D information.



Figure 5. Selected 2D edges for an exemplary window in both images.

However, this process is hindered by the frequent occurrence of such façade structures, which results in a large number of mismatches. Therefore, only these 2D edges are further processed, which are within windows regions already detected within the LIDAR data. Figure 5 depicts selected 2D edges for an exemplary window in both images, which are used to generate 3D primitives by the matching step. Possible mismatches will result in 3D edges outside the available window area. Therefore, these wrong edges can be easily identified and removed. In addition, only horizontal and vertical 3D edges are considered for the further reconstruction process.

For each window, hypotheses about the configuration of the crossbars are generated and tested against the 3D edges derived from the images. Possible shapes are dynamically generated as templates by recursively dividing the window area in two or three parts. Recursion stops when the produced glass panes are too small for a realistic generation of windows. The minimum width and height of the glass panes are restricted by the same threshold value. After each recursion step, the fitting of the template with the 3D edges is evaluated. The partition is accepted if 3D edges are available within a buffer area around the dividing line. In a final step, the crossbars and the window frame are modelled. For this purpose, new 3D cells with a pre-defined thickness are generated at the accepted horizontal and vertical division lines as well as at the window borders. The result is exemplarily shown for two windows in Figure 6.



Figure 6. Reconstructed crossbars for two windows.

In order to stabilize the modelling process of crossbars, neighbourhood relationships are taken into account. The crossbar configuration is assumed to be equal for all windows of similar size which are located in the same row or column. Based on this assumption, similar windows can be simultaneously processed. Thus, the crossbar reconstruction leads to robust results even for windows that are partially occluded or feature strong perspective distortions in the respective image areas.




Figure 7. Refined facade with detailed window structures.

The final result of the building façade reconstruction from terrestrial LIDAR and photogrammetric modelling can be seen in Figure 7. This example demonstrates the successful detection of crossbars for windows of medium size. However, the dynamic generation of templates even allows for the modelling of large window areas as they often occur at facades of big office buildings.

3. FACADE MODELLING FROM LASERMAPS

Our approach based on cell decomposition allows for the explicit refinement of the respective 3D building models. However, efficient visualization of small detail is feasible alternatively, if the 3D representation of the coarse building geometry is separated from the representation of fine surface detail of the facades. This approach is especially suitable, if dense measurements from terrestrial laser scanning are available.

The idea to generate such a combined model of the data, which is especially efficient for rendering purposes was inspired by the concepts for modeling developed in computer graphics. In computer graphics the duality of coarse over-all geometry and fine detail has long been noted. The separation of the two is a fundamental modeling principle. Starting with the observations of Blinn (1978), that the effect of fine surface details on the perceived intensity is “primarily due to their effect on the direction of the surface normal ... rather than their effect on the position of the surface”, modeling concepts were developed, which keep fine surface detail separate as a perturbation of the normal direction or a displacement to the underlying coarser geometry. This idea is a first step in separating the coarse and fine geometry of a surface. While the fine details are not modeled as explicit geometry, their perturbation on the direction of the surface normal is recorded and introduced to the computation of surface intensity. This idea of generating normal maps to enhance the realism of computer generated

surfaces is known as bump mapping and has become a standard in computer graphics. This feature is implemented in hardware on nearly all up-to-date computer graphics card (see for example (Kilgard 2000) and therefore is ideally suited for high performance rendering and real-time animation.

The applicability of this approach for refined visualisation of building model is rather simple. There, the coarse building geometry is described by polyhedral boundary representation with facades typically represented as planar polygons. The exploitation of bump and displacement maps is then similar to the use of intensity images for texture mapping. There, texture coordinates map two-dimensional coordinates of the local planar coordinate system centered on the facade polygon to two-dimensional pixel coordinates of the texture image.

3.1 Generation of LASERMAP

In order to use this 2D-2D mapping for the integration of terrestrial LIDAR data, we have to derive a two-dimensional representation of the point cloud. Obviously this is not possible for the whole point cloud. We rather have to split the point cloud into groups with respect to the facades of the building using a simple buffer operation for each facade polygon. We demonstrate the approach for a case study using the central historical building of the Universitaet Stuttgart. To provide a homogeneous density, the point cloud was acquired with a Leica HDS 3000 from more than 15 stations. The data covers the facades of the building at a point spacing better than 20 mm. Special care was taken to cover all details of the facade and avoid shadowing of occluding objects. For the remainder of this paper we assume that the registration has been performed and the LIDAR data is in the same coordinate frame as the building model.




Figure 8. Existing building model and densely measured 3D point cloud

After this alignment between the building model and the point cloud depicted in Figure 8, the portion of the point cloud that belongs to a particular facade can now be interpolated into a regular raster. This interpolation is very similar to the generation of a Digital Elevation Model from aerial LIDAR. We refer to such a re-interpolated point cloud as a LASERMAP (Böhm, 2005). The term is composed from two terms describing the source of the data, a laser scanner, and the use of the data as a source for 2D mapping.



Figure 9. LASERMAP of a single facade derived from the point cloud.

Figure 9 exemplarily shows such a LASERMAP of the front facade of the aforementioned building. The gray values correspond to offsets relative to the plane of the facade. The map was computed at a resolution of 10 mm to preserve details, which gives an image of 2878 x 1778 pixels. Each pixel stores the offset in 16 bits.

3.2 Normal Map

Since the LASERMAP can be processed like any other image, the generation of a normal map can be implemented easily. This normal map stores the perturbations of the normal vector at each pixel due to surface irregularities. The required unit normal vectors are generated from the partial derivatives of the surface functions as represented by the LASERMAP. These derivatives are computed easily by applying suitable filters, which are well known from image processing.




Figure 10. 3D Rendering of a single LASERMAP as a normal map.

Figure 10 shows the rendering of a single facade polygon using a normal map derived from the surface representation depicted in Figure 9. This example clearly demonstrates that a normal map gives the impression of fine surface detail. However, this is only achieved by varying the shading of each output pixel while the actual geometry is still a flat polygon. This is advantageous as it does not increase the polygon count of the full, when compared to the original model. Still, the flatness of the surface is revealed to the observer under very oblique viewing angles.

3.3 Displacement Map

To generate true 3D geometry is more involved. In order to apply the offset values stored in the LASERMAP, first we have to generate vertices (triangles) which we can apply the offsets to. But instead of explicitly generating vertices and meshing them, we generate the vertices implicitly by using a subdivision

surface (Catmull, 1974). A subdivision surface iteratively subdivides the faces of a given mesh, in our case the facade polygon, into smaller faces. Due to the recursive nature of subdivision surfaces, the number of generated faces is typically a power of two. After the subdivision is defined the LASERMAP can be applied as a displacement map. Figure 11 shows the model of the University building using displacement mapping. For the front facade mentioned above, we created 1024 x 512 faces using subdivision surfaces.



Figure 11. 3D Rendering of the complete model with full displacement maps.

3.4 Automated Repair

It is a common problem in real-world terrestrial laser scanning projects, that scanning artifacts impair the quality of the point cloud. Such artifacts can result from simple occlusions, varying surface reflectivity, beam deflections and other typical problems. It can be very difficult to correct these situations directly in the point cloud, since we can only interact on individual points. Point cloud editing is typically limited to removing points. The situation improves with meshed point clouds, but still interactive editing can be very time consuming.




Figure 12. Semi-automatic repair of an artifact at a half-opened window. top) original 3D rendering, bottom) interactive marking of the defect (white) and the automated match of a replacement (black).




Figure 13. top) result of the automated repair of the LASERMAP, bottom) 3D rendering from the repaired data set.

Since the raster representation of LASERMAPs is essentially the same as an image, suitable operations from image processing can be employed to aid in this task in that case. Thus, the removal of artifacts is considerably simplified and at least partial automation is feasible. In our approach we capitalize on the fact that facade architecture does not consist of purely random geometry, but is composed of repetitive elements. If one of these elements is defect due to artifacts in the point cloud, it can simply be replaced with a copy of an intact element. This repair process is demonstrated in Figure 12 and Figure 13.

As it is already visible in the rendering shown in Figure 10, one of the windows visible in the top row clearly has a defect. This results from the fact that the window was half-opened at the time of scanning. A detailed rendering of this situation is shown again in the top image of Figure 12, while the bottom image shows the corresponding section of the LASERMAP. The repair process is started by interactively marking the defective area in the LASERMAP, shown as a white box in the bottom image of Figure 12. The repair process then automatically searches for a similar area in the LASERMAP. This is done using a simple image matching technique. The best match, which is marked by the black box within this image is then used as a replacement and thus copied onto the defect area.

The result of this repair operation is shown in Figure 13. While the top image gives the corrected LASERMAP, the bottom image shows the corresponding result of the subsequent rendering.

4. CONCLUSION

As it is demonstrated within the paper, the concept of LASERMAPs is a very efficient representation of detailed facade geometry, which is ideally suited both for high-quality rendering and interactive visualization. Bump and displacement mapping allow for a visual improvement of building facades. Additionally, the representation by a LASERMAP enables the development of quick and easy to use tools to remove defects caused by scanning from the facade model. While it is limited

to repetitive structures, it can cope with any kind of shape and is not limited to certain geometries.

However, the resulting model is tied to the resolution of the LIDAR point cloud. When the viewers distance is very close to the surface the limited resolution becomes apparent. To overcome these limitations in resolution, the explicit generation of a 3D CAD model is unavoidable. Additionally, the underlying distance image for the realization of bump and displacement mapping is limited to the representation of 2½ D structures. True 3D structures, which are required for visualization of the highest possible quality, require a modification of the original 3D building model. Within this 3D modeling step also constraints between different object parts like co-planarity or right angles can be guaranteed (Haala et al 2006). Thus, combinations of both approaches will be ideally suited, where CAD primitives are fit to the data where suitable primitives are available and LASERMAPs are used otherwise.

5. ACKNOWLEDGEMENT

We acknowledge the work of Angelos Tsakiroglou, Thanasis Georgantas and Christos Stamatopoulos of the Aristotle University of Thessaloniki, who have acquired the data during their stay in Stuttgart in the ERASMUS exchange program.

REFERENCES

- Catmull, Ed. A. [1974]. Subdivision Algorithm for Computer Display of Curved Surfaces. PhD thesis, Computer Science Department, University of Utah, Salt Lake City, UT, Report UTEC-CSc-74-133.
- Blinn, J. F. [1978]. Simulation of wrinkled surfaces. *ACM SIGGRAPH Computer Graphics* 12(3), pp. 286-292.
- Böhm, J. [2005]. Terrestrial Laser Scanning - A Supplementary Approach for 3D Documentation and Animation, Photogrammetric Week '05, Fritsch (ed.), Wichmann, 263-271.
- Haala, N., Becker, S. & Kada, M. [2006]. Cell Decomposition for the Generation of Building Models at Multiple Scales. IAPRS Vol. XXXVI Part III, Symposium Photogrammetric Computer Vision, pp. 19-24.
- Kilgard, M. [2000]. A Practical and Robust Bump-mapping Technique for Today's GPUs. GDC 2000: Advanced OpenGL Game Development.
- Schuhmacher, S. & Böhm, J. [2005]. Georeferencing of terrestrial laserscanner data for applications in architectural modeling 3D-ARCH 2005: Virtual Reconstruction and Visualization of Complex Architectures, XXXVI, PART 5/W17.