3D Building Generalisation

Martin Kada

Institute for Photogrammetry (ifp)
University of Stuttgart
Geschwister-Scholl-Str. 24D
70174 Stuttgart, Germany
martin.kada@ifp.uni-stuttgart.de

ABSTRACT

The article presents a novel approach for the automatic generalisation of 3D building models with regard to the cartographic visualisation of urban landscapes. Building objects are geometrically simplified for the purpose of levels of detail generation to speed up rendering and to improve the perceptibility on small displays commonly found in mobile devices. Generalisation is realised in a process similar to half space modelling. Approximating planes are determined from the original building model. These are then used as space dividing primitives to remodel the object with a façade and roof structure of simpler shape.

1. INTRODUCTION

The geometric simplification of complex 3D objects has been a topic in various fields of research. In computer graphics, e.g., discrete or continuous levels of detail representations of polygonal or triangular meshes are automatically generated in order to reduce storage space, speed up network transmission, shorten geometric computations and improve rendering performance. During the surface simplification processes, the accuracy of the input model is maintained as good as possible. Due to the geometric error metric, the algorithms show best results on highly complex objects made from hundreds of thousands to millions of primitives. This can be realised even for datasets that are too complex to fit in main memory by out-of-core algorithms (Lindstrom, 2000). A good overview of surface simplification algorithms is given by (Heckbert and Garland, 1997) and (Puppo and Scopigno, 1997).

In contrast to arbitrary objects that can be processed by the aforementioned approaches, single 3D building models are of comparatively low complexity. Collected from data like aerial images or airborne laser scanning, they typically consist of at most a few hundred polygons. Each individual building model, however, but also several objects that make up a building block, exhibits special characteristics that need to be preserved during simplification. These are e.g. right angles that can often be found in the building architecture. A geometric error metric as used in surface simplification does not account for such characteristics. Also because the operators were developed for smooth surfaces rather then for angular 3D shapes, their application to a 3D building model often results in a skewed or tilted model.

Similar problems occur for the automatic simplification of 2D building ground plans. Line simplification, e.g. the algorithm of (Douglas and Peucker, 1973), does not yield very good results for such objects. Consequently, a number of specialised generalisation solutions have been proposed in the past since the early work of (Staufenbiel, 1973). (Sester, 2000) simplifies the shape of 2D ground plans by applying a set of simple rules followed by a least squares adjustment. (Kada, 2002) presents an analogous approach for the generalisation of 3D building models. The polygonal building models are iteratively simplified by combining a number of edge collapse operations into a single step. Building regularities, that need to be detected prior to the simplification, are maintained in the process. Another work on 3D generalisation is based on scale-space theory (Forberg, 2004). Orthogonal building structures are simplified by shifting parallel facets that have been found to be under a certain distance until they merge. A squaring operation is used for non-orthogonal structures such as roofs. (Thiemann and Sester, 2004) presents a segmentation algorithm that is adapted from the work of (Ribelles et al., 2004). The resulting partitioning of the 3D building model is transformed into a Constructive Solid Geometry (CSG) tree. Fragments of the building can then be emphasised, aggregated or eliminated depending on their importance.

All the abovementioned methods are very dependent on the quality of the input models or rely on specific building characteristics. If e.g. angles are close but not exactly orthogonal, building regularities may not be found (Kada, 2002), shifted facets never merge (Forberg, 2004) or the segmentation produces too many fragments (Thiemann and Sester, 2004). Furthermore, the proposed simplification operators are not general enough and only work on a subset of structural elements that can be found in 3D building models.

The paper presents a new generalisation algorithm for 3D building models that overcomes the aforementioned shortcomings. Simplified versions of the respective buildings are generated by remodelling the object using approximating planes. In our context, an approximating plane is the result of averaging a set of polygons that have been identified to belong to the same façade, including protrusions and other small structural elements. In the modelling step, the approximating planes are then used to subdivide an initially infinite space into smaller subspaces. This process is very similar to half space modelling where a cuboid is e.g. modelled by trimming the infinite space using six planar half space primitives. Depending on the definition, the object is then either on the positive or negative side of all six planes. Because our approximating planes have no preferred orientation (see subsection2.1), a recursive subdivision is, however, not possible. Consequently, the algorithm divides all subspaces that intersect a plane by brute force, creating a large number of fragments in the process. These fragments define solids in 3D space that must be further differentiated in building and non-building objects (see subsection 2.2). Non-building fragments are discarded subsequently. Up to this point, the roof polygons have been neglected by the algorithm, leaving flat-top building fragments as an intermediate result. Because the structure of the roof can be very complex, we remodel it individually per fragment using approximating planes averaged from the roof polygons (see subsection 2.3). The resulting objects are referred to as building primitives. The generalised building model is then generated by merging these building primitives.

2. 3D BUILDING GENERALISATION WITH APPROXIMATING PLANES

The presented generalisation algorithm is designed for polyhedral 3D building models. It is assumed that most polygonal wall elements are oriented in parallel, coplanar or rectangular and that these regularities must be preserved during simplification.

2.1. Determining the Set of Approximating Planes

The shape of the generalised building model is highly dependent on the geometric properties of the approximating planes that are used in its generation. Because each subdivision adds complexity to the final model, it is favourable to use the smallest number of planes that optimally approximate the original shape. Each approximating plane represents a set of polygons that belong to the same building façade. These are primarily the polygons with the same orientation. However, also other arbitrary polygons are included that make up protrusions and other structural elements in the façade (see Figure 1.). These elements will be eliminated in the generalised model.

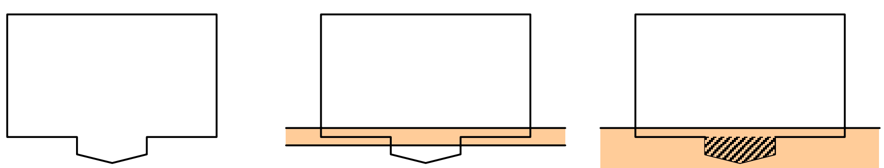

Figure 1. Approximating planes include polygons of the façade and structural elements that are to be eliminated.

In order to keep the number of approximating planes small, polygons with oppositely directed normal vectors are linked to the same plane. The benefit of this approach is that opposing façades under a certain distance threshold will be made coplanar or even merge together (see Figure 2).

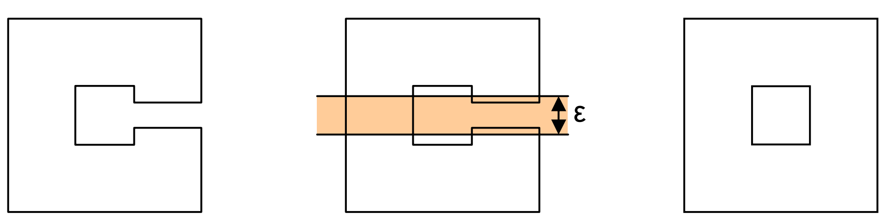

Figure 2. Facades with oppositely directed normal vectors are linked to the same plane and therefore merge together.

A matching of façade polygons to corresponding planes is often ambiguous for complex models even if human interaction is involved. We therefore kept the algorithm for finding the approximating planes rather simple, accepting that a small number of polygons are erroneously linked to planes which represent other façades. This effect can usually be neglected, however, as it only affects small polygons that are not of greater significance on the overall shape of the building.

We use an iterative approach for finding the set of approximating planes. Each iteration results in the plane of most importance, which is measured by the total area of included polygons that are parallel to the approximating plane. At the beginning of an iteration step, a buffer defined by two delimiting, parallel planes is created for each polygon. The initially planar buffers are then merged pair wise to create larger buffers until the delimiting planes reaches a maximum threshold distance. In order to create buffers with many polygons, the orientation of the resulting buffer is averaged before merging. If the polygons from both buffers can not be included in the new averaged buffer, the merge is aborted. The iteration step stops when no more buffers can be merged and the averaged plane equation of the buffer with the largest total area is returned. The polygons inside the buffer are discarded from further processing. By repeating this process, the set of approximating planes is found in descending order of importance. To limit the number of approximating planes, we throw out all planes with a total area that is lower then the square of the threshold distance. See Figure 3(a-f) for example planes determined for a 3D building model.

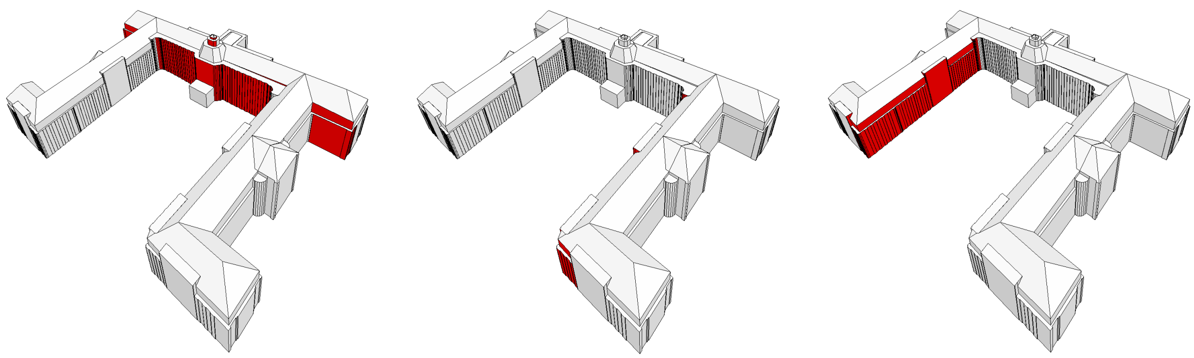

(a) (b) (c)

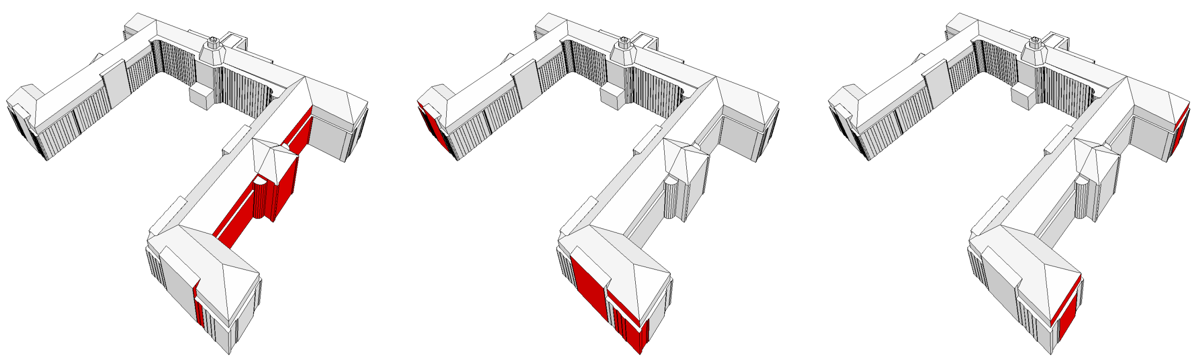

(d) (e) (f)

Figure 3. Thirteen approximating façade planes have been determined for the 3D building model. (a-f) shows six planes and their corresponding polygons (highlighted). (e) Some polygons have been erroneously matched to approximating planes of higher importance (b,c+d).

In order to preserve right angles and also to endorse parallelism in the building model, the approximating planes are analysed in a last step. If the angle of the normal vectors from two or more planes is found to be under a certain threshold, these planes are made rectangular or parallel respectively. If the deviation is only a small angle, this can be done just by changing the normal vector of the plane equation and adjusting the distance value. For larger values, a rotation of the planes around their weighed centroid of the polygons should be chosen.

2.2. Construction of Building Fragments

Once the set of approximating planes have been found, the next step is to generate the building fragments. For this purpose, the first step is to brute force subdivide the infinite 3D space by the set of planes. In practice, an infinite space is unworkable, but a solid two times the size of the building's bounding box is a good substitute for our purpose.

Because the resulting building fragments can not be identified from the subdivision process itself, the objects must be differentiated in building and non-building fragments in a subsequent step (see Figure 4). For that reason, a percentage value is calculated for each fragment that denotes the fraction of building to non-building space. This is basically the volume of the part of the original building model that is located inside the fragment divided by the volume of the fragment itself. All fragments with a percentage value under a given threshold value are then denoted as non-building fragments and are discarded from further processing.

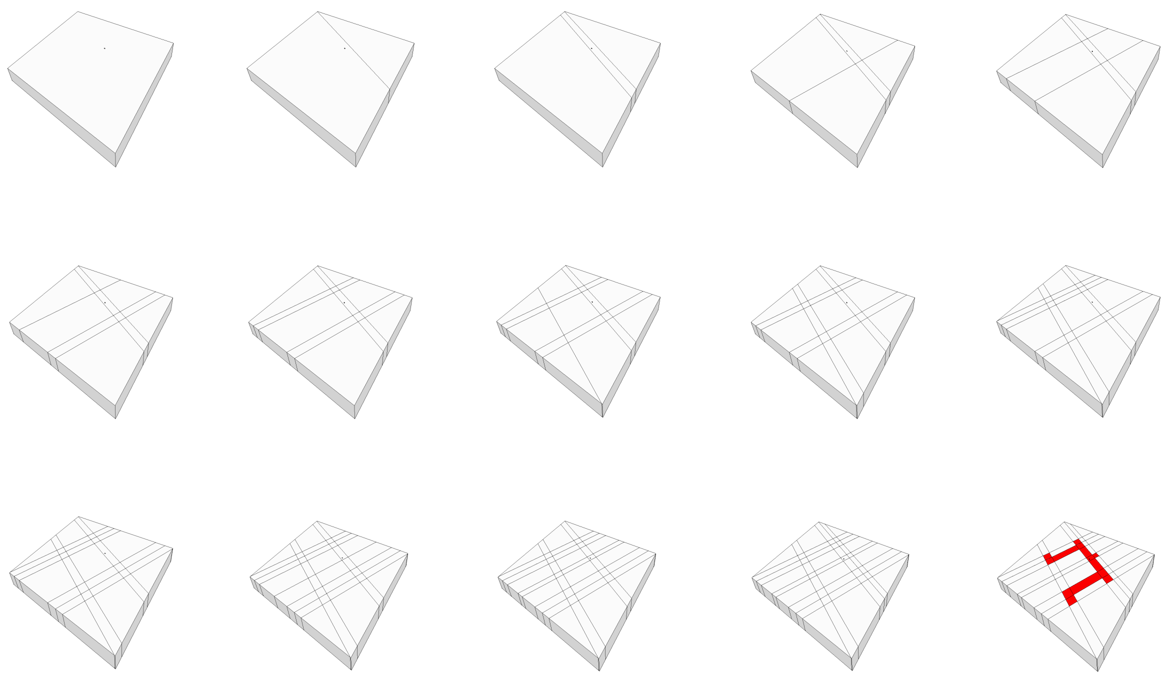

Figure 4. The infinite space is brute force subdivided by the approximating planes. Building fragments are then differentiated in building (highlighted) and non-building fragments.

2.3. Roof Structures

Because the roof structure of 3D building models can be very complex, the roof polygons have been neglected in the generalisation up to this point. Now that the model is fragmented in smaller parts, the generalisation of the roof can be simplified by individually processing the fragments.

However, it has to be ensured that the resulting roof polygons of neighbouring fragments still fit against each other. Therefore, the approximating roof planes are first determined globally from the original building model using the same approach as described in subsection 2.1. Without changing the geometric properties of the planes, a new subset is created for each fragment that is composed of planes that have polygons intersecting the respective fragment. The fragments are then subdivided by their individual set of approximating planes, where the resulting non-building fragments are identified and discarded as described in subsection 2.2 (see Figure 5). The last step is to merge the resulting building primitives into the generalised 3D building model.

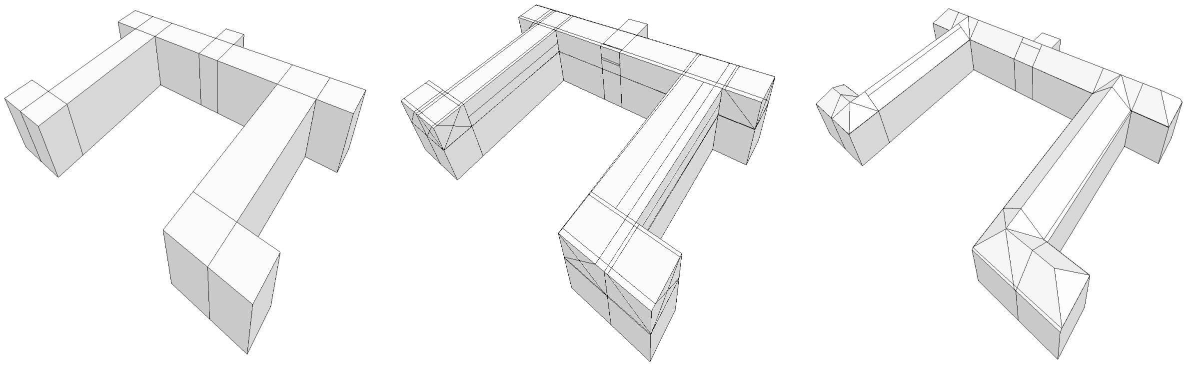

Figure 5. The building fragments are subdivided by the globally determined roof planes and non-building fragments discarded afterwards.

3. IMPLEMENTATION AND RESULTS

The algorithm has been implemented and tested on a number of 3D building models. It shows good results for models ranging from simple to very complex shapes. An assorted selection of example buildings can be seen in Figure 6. The complexity of the objects could be highly reduced in all cases without destroying the overall appearance of the building. In some cases, however, the symmetries could not entirely be preserved because of incorrect symmetries that have been detected while generating the set of approximating planes. Regarding the performance of our implementation, we found out that most computing time was spend on the way we differentiate building and non-building fragments. At this time, a brute force point in polyhedron test is used for a large number of points in a three dimensional regular grid. Still, the computation time for processing our example building models remained under one minute.

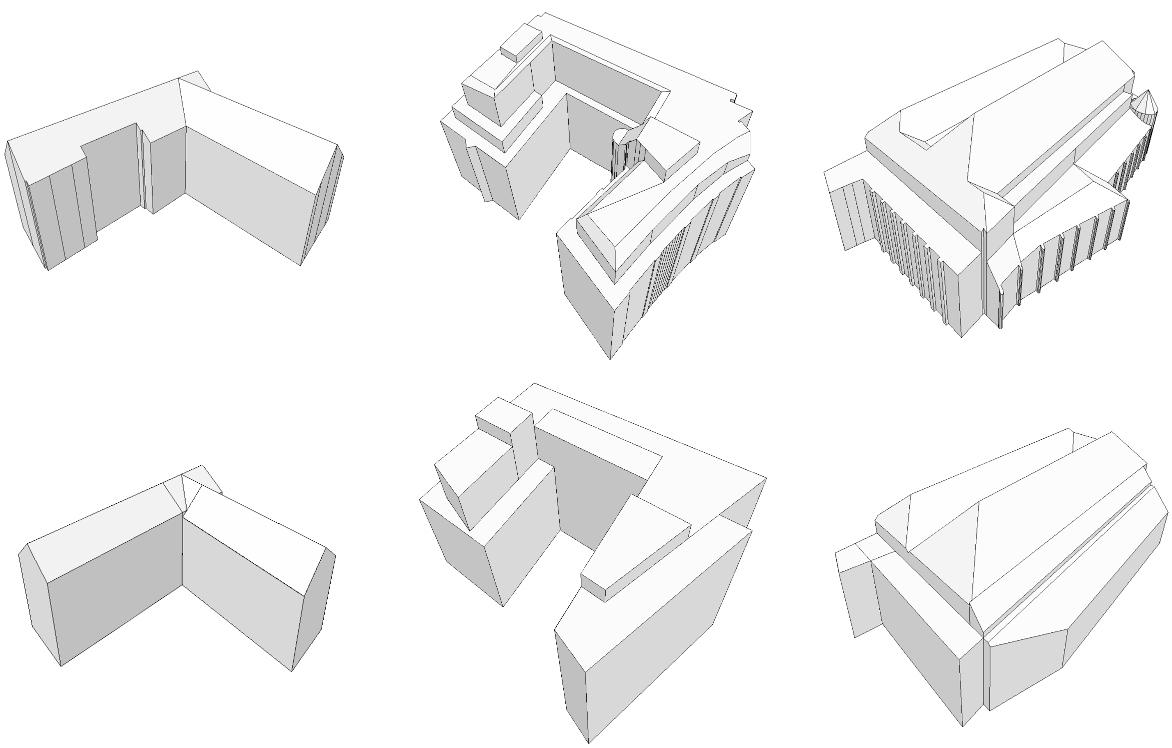

Figure 6. 3D building objects in their original (top) and generalised shape (bottom).

4. CONCLUSION AND OUTLOOK

This article described a new approach for the generalisation of 3D building models. In contrast to other 3D generalisation algorithms that have been proposed in the past, the building model is not iteratively simplified. Instead, the resulting model is generated in a single step by a process similar to half space modelling. By using approximating planes as subdividing primitives, building regularities like coplanarity, parallelism and rectangularity are preserved or can even be enforced if needed. The generalisation process is solely controlled by an intuitive distance threshold that specifies the minimum size of building elements that are created.

So far our focus was on the simplification of single building objects. Each model is generalised individually in disregard of its surrounding area. This can cause problems as the resulting models of neighbouring buildings might overlap. Also the regularities between neighbouring building models may be violated if no special care is taken. Future work will therefore concentrate on the generalisation of several objects in a building block. The resulting models can either be single objects if object identification is required, but also the generation of aggregated objects are feasible.

5. ACKNOWLEDGEMENTS

The research described in this paper is founded by "Deutsche Forschungsgemeinschaft" (DFG – German Research Foundation). It takes place within the Centre of Excellence No. 627 "NEXUS – SPATIAL WORLD MODELS FOR MOBILE CONTEXT AWARE APPLICATIONS" at the University of Stuttgart. The geometry of the 3D building models is provided by Stadtmessungsamt Stuttgart.

6. REFERENCES

Douglas, D. and Peucker, T., 1973. *Algorithms for the reduction of the number of points required to represent a digitized line or its caricature*. The Canadian Cartographer 10 (2), pp. 112-122.

Forberg, A., 2004. *Generalization of 3D Building Data Based on Scale-Space Approach*. In: International Archives of Photogrammetry and Remote Sensing, Vol. XXXV, Part B, Istanbul, Turkey.

Heckbert, P. S. and Garland, M., 1997. *Survey of Polygonal Simplification Algorithms*. In: Multiresolution Surface Modeling Course, SIGGRAPH '97, Los Angeles, USA.

Kada, M, 2002. *Automatic Generalisation of 3D Building Models*. In: Proceedings of the Joint International Symposium on Geospatial Theory, Processing and Applications '02, Ottawa, Canada.

Lindstrom, P., 2000. *Out-of-Core Simplification of Large Polygonal Models*. In: Proceedings of ACM SIGGRAPH 2000, New Orleans, USA.

Puppo, E. and Scopigno, R., 1997. *Simplification, LOD and Multiresolution – Principles and Applications*. In: Eurographics '97 Tutorial Notes, Budapest, Hungary.

Ribelles, J., Heckbert, P. S., Garland, M., Stahovich, T. F. and Srivastava, V., 2001. *Finding and Removing Features from Polyhedra*. In: Proceedings of DETC'01, Pittsburgh, USA.

Sester, M., 2000. *Generalization Based on Least Squares Adjustment*. In: International Archives of Photogrammetry and Remote Sensing, Vol. XXXIII, Part B4, Amsterdam, Netherlands.

Staufenbiel, W., 1973. *Zur Automation der Generalisierung topographischer Karten mit besonderer Berücksichtigung großmaßstäbiger Gebäudedarstellungen*, Ph.D. thesis, Fachrichtung Vermessungswesen, Universität Hannover.

Thiemann, F. and Sester, M., 2004. *Segmentation of Buildings for 3D-Generalisation*. In: Proceedings of the ICA Workshop on Generalisation and Multiple Representation, Leicester, UK.