3D Generalization Lenses
for Interactive Focus + Context Visualization of Virtual City Models

Matthias Trapp, Tassilo Glander, Henrik Buchholz, Jürgen Döllner

Hasso-Plattner-Institute (University of Potsdam)

{matthias.trapp, tassilo.glander, henrik.buchholz, doellner}@hpi.uni-potsdam.de

Abstract

Focus + context visualization facilitates the exploration of complex information spaces. This paper proposes 3D generalization lenses, a new visualization technique for virtual 3D city models that combines different levels of structural abstraction. In an automatic preprocessing step, we derive a generalized representation of a given city model. At runtime, this representation is combined with a full-detail representation within a single view based on one or more 3D lenses of arbitrary shape. Focus areas within lens volumes are shown in full detail while excluding less important details of the surrounding area. Our technique supports simultaneous use of multiple lenses associated with different abstraction levels, can handle overlapping and nested lenses, and provides interactive lens modification.

Keywords—**3D Lenses, Focus + Context Visualization, Virtual 3D City Model, Generalization**

1   Introduction

Today's large-scale virtual 3D city models are characterized by a large number of objects of different types, manifold structures and hierarchies among them, and a high degree of visual detail [5]. Thus, they transport a huge amount of different information, e.g., encoded in facade textures, aerial photographs, building models, infrastructure models, and city furniture. This frequently leads to perceptional and cognitive problems for the user due to visual noise and information overload and, therefore, impairs tasks and usability, e.g., with respect to orientation and navigation in geovirtual environments [9].

To facilitate comprehension, interaction, and exploration of a city model, the cartographic principle of generalization can be applied to create an abstract representation of the city model [9, 16]. Thus, the overall amount of information is reduced while the most important structures are preserved and even highlighted. For large-scale city models, generalized representations at different levels of structural abstraction are needed to achieve appropriate representations at different scales. The abstraction both merges objects that are closely related into higher-level vi-

sualization units and preserves selected landmark objects. The generalization, however, does not provide means for local modifications of the abstraction level, e.g., to integrate more details along a specific route or within a specified area.

The concept of focus + context visualization addresses this need by combining different representations of a model in a single image [25]. For a 3D virtual environment, 3D lenses can be used as a metaphor to control the focus + context visualization [24]. They direct the viewers attention to the focus region, i.e., the model inside the lens, and simultaneously preserve the context information, i.e., the model outside of the lens.
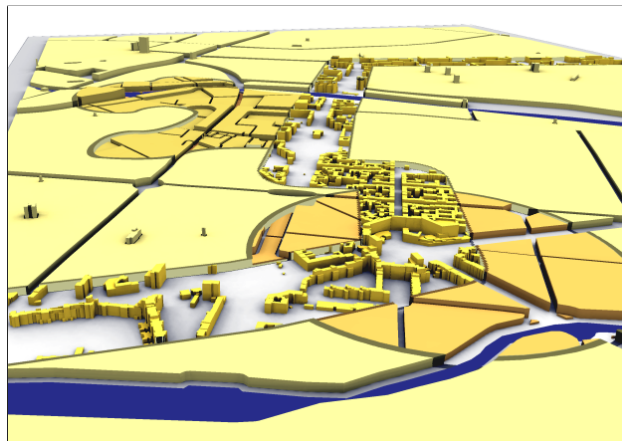

Figure 1: A route through the virtual city model is visualized using our generalization lenses: The route is presented in detail, the context is shown generalized. In addition, two more lenses show different degrees of generalization.

To use 3D lenses most effectively, their interactive manipulation is required, allowing the user to position, rotate, and scale the lens dynamically. When dealing with complex models, the demand for interactivity leads to a number of challenges that have not been solved by existing approaches:

1. **Combination of Multiple 3D Lenses** To achieve a wide scope of possible lens configurations, the vi-

sualization should not be restricted to a single lens. Therefore, the visualization concept and technique have to deal with multiple, intersecting, overlapping and nested 3D lenses that can be freely combined and moved independently. The behavior in these cases should be configurable and consistent.

2. **Arbitrary Lens Shapes** In a given visualization scenario, the core parts to be highlighted cannot be assumed to form a simple shape. Therefore, users should be able to customize the shape of the lenses to fit the actual core parts appropriately.

3. **Data Complexity** The underlying visualization technique must work fast enough to cope with large amounts of geometry and texture data and to support rendering as well as lens manipulation in real-time.

As a matter of principle, existing techniques relying on image-based multi-pass rendering usually have performance problems when applied to large scale 3D scenes with multiple 3D lenses.

Addressing the challenges mentioned above, our work makes the following contributions: We present *Generalization Lenses*, a concept and visualization technique that extends the 3D generalization described in [9] based on *Volumetric Depth Sprites* (VDS) [23] to obtain a customizable visualization, in which the local level of abstraction can be interactively controlled. For this, we introduce a priority based mapping between lens volumes and lens content that handles intersections and nested lenses.

Our interactive visualization technique supports multiple, dynamic, and arbitrarily shaped lens volumes. Figure 1 shows a 3D lens containing a detailed version of a virtual city model integrated into a generalized context.

This paper is structured as follows. Section 2 gives an overview of the related work. Section 3 presents the basic concepts and methods of our approach. Section 4 describes the implementation of our visualization technique and Section 5 discusses the results. Finally, we present ideas for future work (Section 6) and conclude (Section 7).

2 Related Work

Focus + context visualization enables the user to access both, high-level context information and low-level details. Our technique can be seen as a 3D instance of the *Magic Lens*[TM][2, 24] metaphor. We use it to implement focus + context visualization applied to the structural scale [25] of a virtual 3D city model. Only few authors address the application of these magic lenses on 3D geovisualization and virtual 3D city models.

A 2D approach that is comparable with our work can be found in [13]: A generalized (chorematic) and an original version of a 2D map is combined to facilitate users orientation and navigation. The combination and chorematization

cannot be done automatically. Texture lenses, presented in [4], enable the user to combine different texture layers on top of a digital terrain model. This approach is limited to 2D textures and cannot handle overlapping lenses.

2.1 3D Lenses

3D lenses were first introduced in [24], extending the idea of lenses to three dimensional models and scenes.

In [19] an overview of focus + context visualization with 3D lenses and application examples is given. Based on [18], they present an image-based multi-pass algorithm to separate focus from context regions. The algorithm supports arbitrarily shaped lenses in real-time, but does not handle overlapping or nested 3D lenses.

In [23], we presented a volumetric test that can be used for focus and context separation. Using depth peeling [3], the approach transforms arbitrarily shaped solids (convex and non-convex) into layered depth images [21] and performs ray-casting in 3D texture space to determine the parity of a 3D point to test. This test can be applied at various levels (vertex, primitive, fragment) within the rendering pipeline.

2.2 Generalization of 3D City Models

In cartography, the term *generalization* means to abstract meaningful. It describes the process of reducing details of the depicted spatial information to a degree that is appropriate to scale, task and viewer [10]. Due to the subjective nature of this process, designing an automatic algorithm is a challenge. For the 2D case, solutions have been found, e.g., based on agents [6] or based on least squares adjustment [20].

For 3D building generalization, existing approaches concentrate on the simplification of single buildings. For this, they remodel a building with a set of characteristic planes [11], or split it along characteristic planes into a CSG tree representation [22]. Morphological operations have also been applied to building generalization [7]. In [17], the generalization depends on line simplification performed on the projected walls [1].

The approaches described above are limited to single objects and disregard aggregation of multiple objects. In an earlier work, we introduced a 3D generalization technique that performs aggregation of multiple objects [8, 9]. In this work, this technique is used as a preprocessing step.

3 Concept

The data processing steps in our concept can be grouped into a preprocessing phase and a rendering phase (Figure 1).

Preprocessing Phase This step prepares all necessary geometry for the rendering phase. This includes the generalization of city structures (Section 3.1) into a sequence ($\mathcal{LOA}$) of different levels of abstraction,
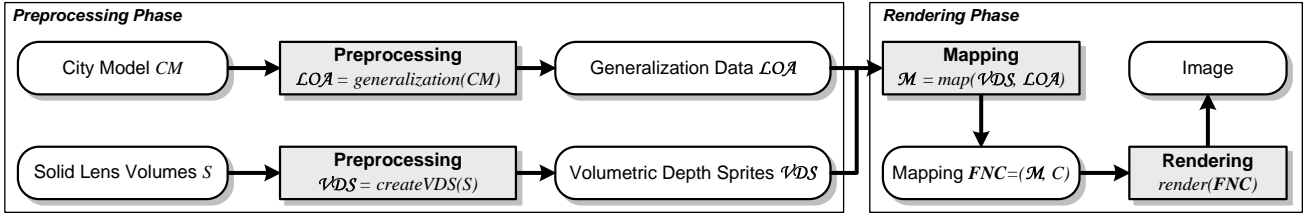
Figure 2: Visualization pipeline of our concept. A preprocessing phase prepares 3D generalization geometry and the 3D lens volumes for the rendering phase. The mapping of both can be specified dynamically at runtime.

the creation of volumetric depth sprites (Section 3.2) to represent the lens volumes ($\mathcal{VDS}$), and an initial mapping $\mathcal{M}$ between both (Section 3.4).

Rendering Phase During runtime, the mapping $\mathcal{M}$ between levels of generalized geometry $\mathcal{LOA}$ and 3D volumes ($\mathcal{VDS}$) of the lenses can be specified and modified by the user. The complete focus + context mapping FNC is rendered in subsequent passes: In each pass we apply pixel-precise clipping against the lens volumes (Section 4.2). Every LOA geometry is rendered only once per frame.

3.1 Discrete Levels of Generalizations

The generalization technique creates a sequence of city model representations with increasing levels of abstraction. The generalization representation LOA_i of level i generalizes LOA_{i-1}. More specifically, one component from LOA_i aggregates a number of components from LOA_{i-1}. As our generalization technique depends on weights given to the infrastructure, the number of generalization representations is the number of weight classes plus one.

Our technique focuses on aggregation, as it implicates the strongest abstraction compared to other generalization operators such as simplification. The number of single objects is significantly decreased when turning to the next level of abstraction. For example, the city model used in this paper contains $10,386$ objects in LOA_0. It is reduced to 468 objects in LOA_1 and to 66 objects in LOA_7. To obtain a homogeneous visualization, no facade textures are kept in the generalized representations. Textures are only used to provide depth cues through a lighting texture [5].

3.2 Volumetric Depth Sprites

To separate focus from context with per-pixel precision, we use VDSs and a *Volumetric Parity Test* (VPT) introduced in [23]. The concept allows us to efficiently perform clipping against multiple, arbitrarily shaped volumes within a single rendering pass. Depending on the *parity* of a VDS, the rendering algorithm clips either the geometry within the represented volume or its complement. The *active* state indicates if this operation is performed.

3.3 3D Lens Shapes

At runtime, the lens shapes are stored as VDS representations, and therefore, can be scaled, rotated and translated within the scene. Our system supports the following methods for the lens shape creation:

Derived Shapes Our framework can generate lens shapes from buffered 2D polygonal shapes and polylines. This allows us to derive complex lens shapes directly from geo-referenced data.

Modeled Shapes Lens shapes can also be modeled explicitly using 3D modeling software by importing these through common interchange formats.

3.4 Mapping Lenses to Generalization Levels

The mapping $\mathcal{M} = \{M_i | i = 0 \ldots n\}$ between lens shapes $VDS_i \in \mathcal{VDS}$ and generalization levels $LOA_i \in \mathcal{LOA}$ can be described as a tuple:

$$M_i := (VDS_i, LOA_i) \tag{1}$$

Here, i denotes the priority of the mapping. This explicit order upon generalization levels is necessary to handle overlapping volumes correctly. VDS_i represents the focus volume whose content is defined by the generalization level LOA_i. A particular generalization level represents the context C. The complete mapping in terms of focus + context visualization is defined as:

$$FNC := (\mathcal{M}, C) \tag{2}$$

4 Implementation

Our approach has been implemented based on OpenGL [15] in combination with GLSL [12] for rendering. We use the Computational Geometry Algorithms Library (CGAL) to calculate the polygonal cells used during the generalization process [26].

4.1 Generalization Algorithm

The generalization is done as a preprocessing step. Input from the city model are the infrastructure network and the building geometry. The infrastructure network consists of a number of weighted polylines, as provided by commercial (e.g. TeleAtlas [R]) or non-commercial (e.g. OpenStreetMap) geo-data providers. The building models can

be separated into a large number of prismatic block buildings and a small number of high detail CAD-based buildings. Typically, important landmark buildings are modeled in high detail, while the majority of buildings are relatively simple, as they are automatically created using remote sensing technologies.

The preprocessing starts with the creation of polygonal cells by computing the arrangement of the polylines. Using cells defined by the streets can be argued with the cognitive importance of streets to structure the city [14]. Then, the building geometry is mapped to the cells using point-in-polygon tests. Then, for each cell a new height is determined by calculating the mean height of all buildings within one cell, thereby respecting the proportionate building footprint's area.

Important landmark buildings that are essential for orientation are preserved during block creation. We select buildings, of which the height surpasses the cell's mean height $\overline{h}$ by twice the standard deviation σ, that is, if a building b fulfills this condition:

$$height(b) > \overline{h} + 2 \cdot \sigma \qquad (3)$$

If a data set provides additional hints, these can also be used for landmark selection, e.g., for block models that are enriched by detailed 3D models for selected buildings. In addition, the set of preserved landmark buildings can be explicitly adjusted.

4.2 Rendering of 3D Lenses

To render the mapping FNC, our technique performs multi-pass rendering with one pass per level of abstraction.

Figure 3 shows the pseudo code for rendering the mapping described in Section 3.4. Starting with rendering the context geometry C, our algorithm performs pixel-precise clipping against all VDS_i within a single rendering pass. After that, we render the geometry of the generalization layers successively. Beginning with the lowest priority level, we swap the parity of the associated VDS,

```
render(FNC){
  ∀ Mᵢ ∈ M {
    VDS ← VDSᵢ ∈ Mᵢ
    setActive(VDS, true)
    setParity(VDS, false)
  }
  renderGeometry(C)
  i = ∥ M ∥
  while(i > 0){
    VDS ← VDSᵢ ∈ Mᵢ
    LOA ← LOAᵢ ∈ Mᵢ
    if(!culling(VDS){
      setParity(VDS, true)
      renderGeometry(LOA)
      setActive(VDS, false)}
    i = i-1
}}
```

Figure 3: Pseudo code for rendering multiple levels of generalizations.

render the geometry LOA, and turn off the current VDS. This ensures that the geometry of lower priority does not interfere with geometry of higher levels. This algorithm is easy to implement and exploits the design of the VPT.

To increase runtime performance, we apply view-frustum culling to the 3D lenses. If no corner vertex of the transformed bounding box of a VDS is inside the current view frustum, we set the active status of the respective VDS to false.

5 Results & Discussion

Figure 1 shows an application of our approach for the interactive visualization of large scale virtual 3D city models. The used input dataset comprises the inner city of Berlin with about $10,386$ generically textured buildings on top of a digital terrain model.
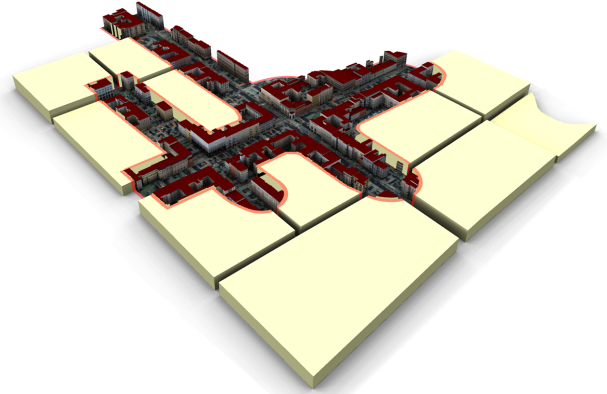

Figure 5: Application example for a single, arbitrarily shaped 3D lens. The focus preserves important details along a specified route. The context region is generalized.

5.1 Usage Scenarios

We have tested our approach in different usage scenarios.

Single Focus is the standard use case for our lens visualization where the mapping is usually defined as:

$$FNC = (\{(VDS_0, LOA_i)\}, LOA_j), j > i \qquad (4)$$

Figure 5 shows an example. It emphasizes a single region of interest.

Multiple Foci implicate the handling of disjunctive, overlapping, and nested regions of interest. Figure 4.A and B show examples of two overlapping regions of interest.

We implemented two types of lenses. As described in [19] they can distinguished by the modification of the lens position during rendering with respect to the camera. Our visualization technique supports a mixture of both lens types:

Scene Lens The position of this lens type is independent from the user's orientation. The lens position can be fixed with respect to the virtual environment or attached to a moving object in the scene.

Figure 4: Application examples for 3D Generalization Lenses. Figure A and B show multiple overlapping lenses with different mappings. Figure C shows two nested, camera lenses.

Camera Lens This lens type adapts it position with respect to the current user orientation. It can be used to assure that potential foci are always visible (Figure 4.C). This minimizes the effort for the user to steer the lenses.

5.2 Performance

Our test platform is an NVIDIA GeForce 8800 GTS with 640 MB video memory and Athlon™64 X2 Dual Core 4200+ with 2.21 GHz and 2 GB of main memory at a viewport resolution of 1600x1200 pixel. The test application does not utilize the second CPU core.

All scenes depicted in this paper can be rendered at interactive frame rates. The rendering performance depends on the number and depth complexity of the VDS used, hence, from the number of samples the VPT has to perform. Further, it is limited by the geometrical complexity of the generalization levels.

Pixel-precise clipping against multiple VDS is fill-rate bound. Consequently, the performance is proportional to the distance between user and rendered geometry, i.e., it improves with increasing distance to the lens.

5.3 Limitations

The main conceptual limitation refers to the one-to-one mapping between generalizations levels and lens volumes. Therefore, it is currently not possible to assign multiple volumes to a single generalization level.

Further, our concept is limited by the memory consumptions of the *LOA* and *VDS*. The storage of high quality generalization levels, e.g., with precomputed lighting textures, exceeds easily the main memory size.

The main drawback concerns the rendering performance that depends on the number and depth complexity of the VDS. To reduce visual artifacts during clipping we colorize the visible back faces. Therefore it is not possible to apply back-face culling.

6 Future Work

Currently, we are extending our implementation to eliminate rendering artifacts such as shading discontinuities for visible back faces. We also research a loss-less compression algorithm for volumetric depth sprites to minimize video memory consumptions and to optimize the number of necessary texel fetches for the VPT. In addition, we want to extend the mapping to facilitate the binding of multiple lens volumes to multiple generalization levels.

To further demonstrate the flexibility of our VDS approach, we want to apply it to more usage scenarios such as the visualization of statistic geo-referenced data and spatio-temporal data. Focus regions can then be highlighted further by using different rendering styles such as non-photorealistic rendering.

As city models are usually organized in a plane, generalization lenses do not exhaust the potential of the VDS approach. Therefore, we plan to implement in-house lenses to give insight to complex structured interior space, e.g., to visualize escape routes.

7 Conclusions

We present an interactive focus + context visualization technique that combines different levels of generalization of a virtual 3D city model within a single image. Our approach is based on an automated generalization algorithm for virtual 3D city models and pixel-precise clipping against multiple, arbitrarily shaped, polygonal meshes, which act as 3D lenses. The presented 3D generalization lenses enable the implementation of interactive tools that allow users to flexibly combine different model LOAs within the same visualization.

Acknowledgments

References

[1] K.-H. Anders. Level of Detail Generation of 3D Building Groups by Aggregation and Typification. *Proceedings of 22nd International Cartographic Conference, La Corua, Spain*, 2005.

[2] Eric A. Bier, Maureen C. Stone, Ken Pier, Ken Fishkin, Thomas Baudel, Matt Conway, William Buxton, and Tony DeRose. Toolglass and Magic Lenses: The see-Through Interface. In *CHI '94*, pages 445–446, New York, NY, USA, 1994. ACM Press.

[3] Cass Everitt. Interactive Order-Independent Transparency. Technical report, NVIDIA Corporation, 2001.

[4] Jürgen Döllner and Konstantin Baumann. Geländetexturen als Mittel für die Präsentation, Exploration und Analyse komplexer räumlicher Informationen in 3D-GIS. In V. Coors A. Zipf, editor, *3D-Geoinformationssysteme*, pages 217–230. Wichmann Verlag, 2005.

[5] Jürgen Döllner, Henrik Buchholz, and Haik Lorenz. Ambient Occlusion - ein Schritt zur realistischen Beleuchtung von 3D-Stadtmodellen. *GIS - Zeitschrift für Geoinformatik*, pages 7–13, November 2006.

[6] C. Duchêne. Coordinative Agents for Automated Generalisation of Rural Areas. *Proceedings 5th Workshop on Progress in Automated Map Generalization*, 2003.

[7] A. Forberg. Generalization of 3D Building Data based on a Scale Space Approach. *International Archives of Photogrammetry*, 35, 2004.

[8] Tassilo Glander and Jürgen Döllner. Cell-based Generalization of 3D Building Groups with Outlier Management. In *GIS '07: Proceedings of the 15th annual ACMGIS*, pages 1–4, New York, NY, USA, 2007. ACM.

[9] Tassilo Glander and Jürgen Döllner. Techniques for Generalizing Building Geometry of Complex Virtual 3D City Models. In *2nd International Workshop on 3D Geo-Information: Requirements, Acquisition, Modelling, Analysis, Visualisation*, Delft, Netherlands, December 2007.

[10] G. Hake, D. Grünreich, and Liqiu Meng. *Kartographie*. Walter de Gruyter, Berlin, New York, 8 edition, 2002.

[11] M. Kada. 3D Building Generalisation. *Proceedings of 22nd International Cartographic Conference, La Coruña, Spain*, 2005.

[12] John Kessenich. *The OpenGL Shading Language Language Version: 1.20 Document Revision: 8*, September 2006.

[13] A. Klippel and K. Richter. Chorematic focus maps, 2004.

[14] K. Lynch. *The Image of the City*. MIT Press, 1960.

[15] Mark J. Kilgard. NVIDIA OpenGL Extension Specifications. Technical report, NVIDIA, November 2006.

[16] L. Meng and Forberg A. 3D Building Generalization. In W. Mackaness, A. Ruas, and T. Sarjakoski, editors, *Challenges in the Portrayal of Geographic Information: Issues of Generalisation and Multi Scale Representation*, pages 211–232. 2006.

[17] Jiann-Yeou Rau, Liang-Chien Chen, Fuan Tsai, Kuo-Hsin Hsiao, and Wei-Chen Hsu. LOD Generation for 3D Polyhedral Building Model. In *Advances in Image and Video Technology*, pages 44–53, Berlin Heidelberg New York, 2006. Springer-Verlag.

[18] Timo Ropinski and Klaus Hinrichs. An Image-Based Algorithm for Interactive Rendering of 3D Magic Lenses. Technical Report 03/04 - I, FB 10, Institut für Informatik, Westfälische Wilhelms-Universität Münster, 2004.

[19] Timo Ropinski, Klaus H. Hinrichs, and Frank Steinicke. A Solution for the Focus and Context Problem in Geo-Virtual Environments. In *Proceedings of the 4th ISPRS Workshop on Dynamic and Multi-dimensional GIS (DMGIS05)*, pages 144–149, 2005.

[20] M. Sester. Generalization Based on Least Squares Adjustment. *International Archives of Photogrammetry and Remote Sensing*, 33:931–938, 2000.

[21] Jonathan Shade, Steven Gortler, Li wei He, and Richard Szeliski. Layered Depth Images. In *SIGGRAPH '98*, pages 231–242, New York, NY, USA, 1998. ACM.

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

[23] Matthias Trapp and Jürgen Döllner. Real-Time Volumetric Tests Using Layered Depth Images. In E. Reinhard K. Mania, editor, *Proceedings of Eurographics 2008*. Eurographics, The Eurographics Association, April 2008. to appear.

[24] John Viega, Matthew J. Conway, George Williams, and Randy Pausch. 3D Magic Lenses. In *UIST '96*, pages 51–58, New York, NY, USA, 1996. ACM Press.

[25] Colin Ware. *Information Visualization: Perception For Design*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2000.

[26] Ron Wein, Efi Fogel, Baruch Zukerman, and Dan Halperin. 2D Arrangements. In CGAL Editorial Board, editor, *CGAL-3.2 User and Reference Manual*. 2006.