Enhancing the Visibility of Labels in 3D Navigation Maps

Mikael Vaaraniemi, Martin Freidank, Rüdiger Westermann

Abstract The visibility of relevant labels in automotive navigation systems is critical for orientation in unknown environments However, labels can quickly become occluded, e.g. road names might be hidden by 3D-buildings, and consequently, the visual association between a label and its referencing feature is lost. In this paper we introduce five concepts which guarantee the visibility of occluded labels in 3D navigation maps. Based on the findings of a pre-study, we have determined and implemented the two most promising approaches. The first approach uses a transparent aura to let the label shine through occluding objects. The second method lets the feature, e.g. the roads, glow through the 3D environment, thus re-establishing the visual association. Both methods leave the 3D world intact, preserve visual association, retain the label's readability, and run at interactive rates. A concluding user study validates our approaches for automotive navigation. Compared to our baseline – simply drawing labels over occluding objects – both approaches perform significantly better.

1 Introduction

Automotive navigation devices started appearing in the mid-80s. The first commercially available device, the Etak Navigator introduced in 1986, guided drivers with an annotated 2D map and guidance arrows to their destination [29]. Since then, textual annotations in maps have been helping the driver navigate through unknown environments. They are essential for the exploration of navigation maps. The visu-

Mikael Vaaraniemi
BMW Research and Technology GmbH, Munich, e-mail: mikael.vaaraniemi@bmw.de

Martin Freidank
University of Koblenz-Landau, e-mail: mfreidank@uni-koblenz.de

Rüdiger Westermann
Technische Universität München, e-mail: westermann@tum.de

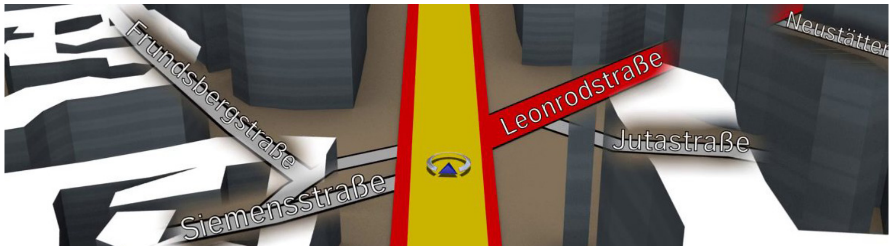

(a) Transparency label aura: the labels blend out occluding 3D objects.

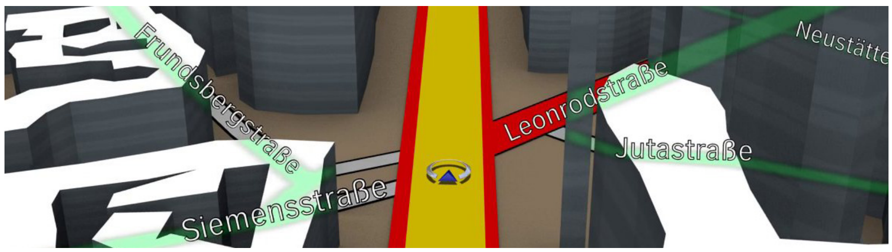

(b) Glowing roads: the roads shine through occluding 3D objects.

Fig. 1 The two selected approaches to preserve the visibility of textual labels in a 3D world.

alization has improved gradually and nowadays, 3D navigation maps have become omnipresent. Several competing companies, like Sygic or Navigon, include terrain and 3D city models in their latest navigation devices. In these systems, labels are usually rendered over occluding 3D elements (e.g. road names over buildings). This approach makes them easily readable, but the visual association to their corresponding feature is lost. As labels appear in front of occluding objects, depth perception is hindered and spatial orientation becomes difficult. In this paper, our primary goal is to preserve the visibility of labels in 3D navigation maps. Hence, deduced from cartographic rules by Imhof [16] and our expert study from section 4, we define the following rules for labeling 3D navigation maps:

- All labels should be readable, even occluded labels
- The visual association between the label and its feature should be guaranteed
- Labels should not occlude other labels or important features
- Depth cues of the 3D world should be preserved
- Labels should support spatial orientation

Our main contribution are two approaches fulfilling these rules and, consequently, enhancing the visibility of occluded labels in 3D navigation maps. The first approach creates a transparency aura around every label and lets labels shine through occluding objects (see Fig. 1(a)). The second method lets the referenced features, e.g. the roads, glow through the 3D environment, thus creating a visual association (see Fig. 1(b)). Both methods leave the 3D world intact, preserve visual association and retain the labels' readability. Also, they are able to run at interactive framerates. The enhancements of these approaches are validated in a user study.

2 Labeling Techniques

2.1 World-Space and Screen-Space Labels

Annotations can be placed in World-Space (WS) or in Screen-Space (SS) into the 3D world. SS labels (or 2D labels) are placed parallel to the screen (see Fig. 2(a), 2(b)). They can be thought as being part of a Head-Up-Display (HUD), overlaid over the 3D scene. WS labels (or 3D labels) are part of the 3D world (see Fig. 2(c), 2(d)). As such, they are transformed by the perspective projection. Chen et al. [4] compare both types of labels. They show that SS labels are better for naive search tasks in densely packed scenes. Also, they are easy to read because they are always facing the viewer. In contrast, as WS labels are part of the 3D scene, they exhibit occlusion problems and can be very difficult to read (e.g. when they follow the object's curvature). However, because they provide strong association cues, they improve the visual association to the referenced feature [9]. Polys et al. [25] evaluate both techniques and state, that even tough WS provides tight coupling, SS performs better across all tested tasks.

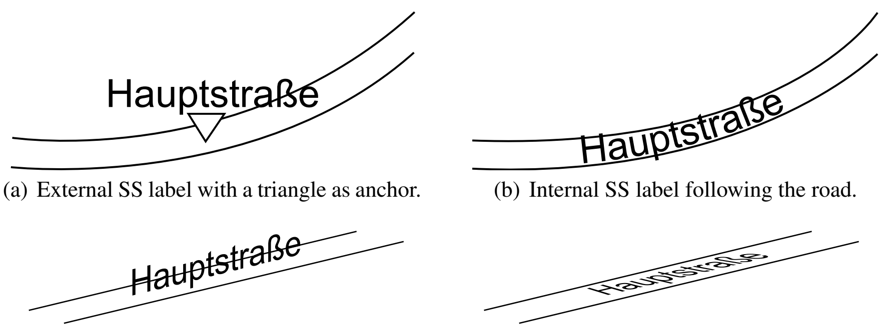

(a) External SS label with a triangle as anchor.

(b) Internal SS label following the road.

(c) Internal WS label placed upright.

(d) Internal WS label laid onto the road.

Fig. 2 World-Space (WS) and Screen-Space (SS) labeling techniques used in our approaches.

2.2 External and Internal Labels in 3D worlds

External Labels. Fekete and Plaisant [8] introduce external labels to annotate dense sets of points. Connected with an anchor (e.g. a line or a triangle), they are displayed beside (or outside) the referenced objects (see Fig. 2(a)). Hence, they do not hide the referenced object. Because they are primarily displayed as SS labels they are also easy to read. External labels are mainly used for annotation of single 3D objects, e.g. in scientific illustrations [13, 1]. However, Maass and Döllner [21] use external labels to annotate virtual landscapes. Their approach creates dense clusters of labels and long connecting lines which makes visual association nearly impossible. Stein et al. [27] compute the placement of external SS labels in a 3D world with an optimization algorithm. To determine the visibility of a label, a sphere is placed at the 3D position of the anchor. Its percentage of occlusion determines the transparency of the label. If the sphere is fully occluded, the feature is not labeled. All these approaches use greedy algorithms to compute an optimum placement for annotations. The computed positions are connected with the referenced object with an anchor line. This connection makes the visual association more difficult compared to a placement directly beside the object. Additionally, as shown by Maass et al. [24], using anchor lines might impair depth perception.

Internal Labels. Internal labels are spatially bound to an object. This allows for a direct visual association to the referenced object (see Fig. 2(b)). For instance, Maass and Doellner [20] annotate 3D buildings intuitively with billboards in WS. They introduce an approach to annotate line features in WS [22]. They determine the placement of labels on the fly using sample points. But, changing the view results in different label placements and thus in a temporally incoherent layout. They present an approach to integrate labels directly onto the hulls of 3D buildings by taking their shape into account [23]. This creates internal WS labels which are part of the world. In general, internal labels depict the visual extent of an object. Ropinski et al. [26] and Cipriano and Gleicher [5] introduce internal WS labels to annotate e.g. medical illustrations. However, these labels hide parts of the referenced object and their readability depends on distortion and the viewing angle.

Hybrids. Bell et al. [2] and Götzelmann et al. [11, 12] present similar hybrid approaches, which use internal and external labels. Bell et al. annotate virtual 3D cities while Götzelmann et al. annotate scientific illustrations. External labels with anchor lines are used when the viewer is far away. When the viewer gets closer and the objects' dimensions allow it, they use internal labels. In contrast, Google Earth [10] uses SS external labels for cities and WS internal labels for streets. This makes street names difficult to read at low viewing angles.

2.3 Summary

None of the presented approaches satisfy our stated goals in section 1. In particular, the goal to preserve readability of labels which are being occluded in a 3D world. The computations of most SS layouting algorithms are done solely in screen space. They do not take into account the occlusion between labels and a 3D scene. Many examples can be found in the bibliography by Wolff and Strijk [31]. Furthermore, SS approaches to annotate scientific illustrations place external labels around single objects, hence, are not affected by occlusion problems [13, 1, 11, 12]. Most SS approaches for labeling 3D worlds ignore occlusion problem by rendering labels over the scene (similar to a HUD) [21, 10]. Only newer SS algorithms take the visibility of the anchor into account [27]. On the other hand, internal WS approaches try to find visible positions for labels at runtime [20, 22, 23]. However, if unsuccessful, the object remains unlabeled.

3 Concepts

In this section we introduce several concepts which assure the visibility and thus preserve the readability of labels occluded by objects of the 3D world.

3.1 Baseline

The first concept we introduce represents our baseline. It consists of drawing the labels over the 3D world (see Fig. 3). Hence, all occlusion created by objects from the 3D world is ignored. We chose it as a baseline, because it is a straightforward solution for resolving occlusion problems. Also, it is used in almost all existing navigation systems, e.g. Sygic GPS Navigation [28] and Google Earth [10].

Fig. 3 Baseline: drawing labels over the 3D world in bird's eye with SS (left) and snail view with WS labeling (right).

3.2 Cutaways

Our second concept is cutaways (see Fig. 4). This method is inspired by 2D magic lenses which were first introduced by Bier et al. [3]. These lenses highlight focus regions by modifying their representation. One such approach Bier et al. depicts, is the wireframe representation inside the focus region. Viega et al. [30] extend these to 3D environments with flat and volumetric lenses. Coffin and Höllerer [6] introduce perspective cutaways for 3D scenes. The resulting holes are rendered with the correct perspective as if they were cut in the occluding object. Our approach is very similar to the perspective cutaways. Every label creates a focus region which cuts away all occluding objects in a perspectively correct manner.

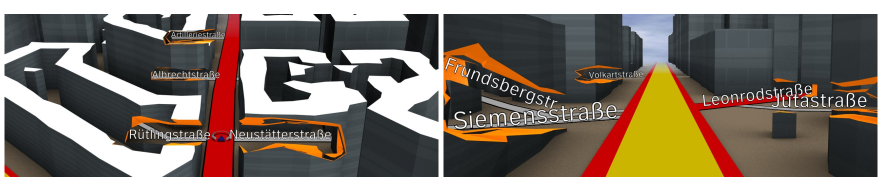

Fig. 4 Cutaways: labels create perspective cut aways in occluding objects of the 3D world in bird's eye with WS (left) and snail's view with SS labeling (right).

3.3 Transparency Label Aura

The next concept creates a smoothly blended transparency aura around the labels. It is similar to Krüger et al. [19] interactive focus+context method called ClearView. Their approach is directly inspired by magic lenses. They create a semi-transparent area around the focus region while the remaining parts stay opaque to preserve context information. Elmqvist et al. [7] evaluate such x-ray vision and state that it leads to faster and better object discovery. Analogously, we define in our concept a transparency region around the label (similar to a focus area). All objects of the 3D world lying in front of this region become transparent. This x-ray vision lets the user read every label. Because we define the region to be larger than the label, the referenced feature (e.g. the road) can be seen partially. This preserves the context of the focus region. Hence, the visual association to the referenced feature is retained.

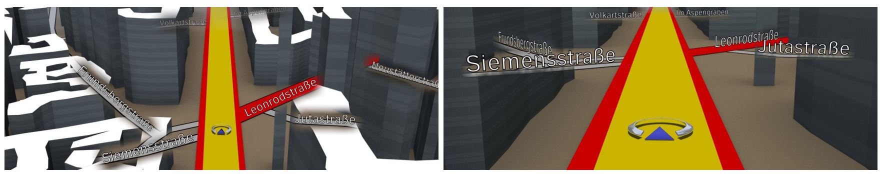

Fig. 5 Transparency label aura: labels create a transparent region in the occluding objects in bird's eye with SS (left) and snail's view with WS labeling (right).

3.4 Glowing Labels

In our third concept we let labels glow through occluding objects (see Fig. 6). This method is inspired by augmented reality (AR) applications. Kalkofen et al. [17, 18] present an approach to augment real objects with context+focus information. This helps recreate the spatial relationship between reality and virtual information. We note that this approach is used in almost all isometric strategy PC games (e.g. Command & Conquer, Age of Empires). Units being hidden by structures (e.g. buildings) are usually tinted with a different color. Similarly, we tint the occluded parts of labels with a color distinct from the surrounding world.

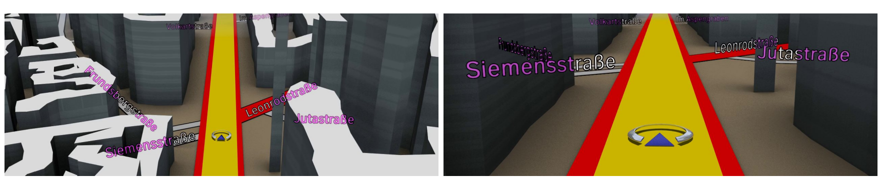

Fig. 6 Glowing labels: labels are glowing through the 3D world with a distinct color in bird's eye with SS (left) and snail's view with WS labeling (right).

3.5 Glowing Roads

The baseline concept makes the labels visible but thereby loses the visual association to its referenced feature, e.g. the road. Our fourth concept tries to solve this problem by adding glowing roads to the baseline. Again, in a similar fashion to the approaches by Kalkofen et al., we let the occluded parts of the roads shine through the 3D world (see Fig. 7). This method recreates the missing context of the labels.

Fig. 7 Glowing roads: roads are glowing through the 3D world in bird's eye with SS (left) and snail's view with WS labeling (right).

4 Expert Study

We conducted an initial expert study. Our goal was to determine which of the introduced concepts fulfills our rules for labeling a 3D navigation map (stated in section 1). Also, we wanted to form an opinion about the usability and aesthetics of each method from our domain experts. Besides, the preferred labeling space (SS or WS) was surveyed. Two engineers working for over five years on automotive navigation were chosen as experts. Also, as further subjects, we selected three research engineers working on human machine interaction systems.

4.1 Study design

We presented the four concepts introduced in section 3: cutaways (see Fig. 4), transparent label aura (see Fig. 5), glowing labels (see Fig. 6) and glowing streets (see Fig. 7). Each concept was compared to our baseline: rendering labels over the 3D scene (see Fig. 3).

Movies. Movement is an important aspect which greatly affects the way a 3D concept is perceived. Animation can cause occlusion and creates an important depth cue: the motion parallax. Hence, to improve the value of our study, we chose to create animated sequences lasting 20 to 30 seconds. Each movie was shown with SS- and WS-labeling. We presented each movie with the same flight path in two perspectives: a snail view closer to the ground and a bird's eye view. All these combinations culminated to sixteen different animated sequences. To each subject we showed these concepts in a fixed order as they are introduced in section 3. In an ensuing discussion, we queried all statements and asked for a ranking of the presented concepts (see Fig. 8).

Conceptual Details. We selected a light violet color for the glowing labels (see Fig. 6). Usually, such a color is not present in a 3D navigation visualization, yet it still remains an aesthetically pleasing color. The hidden parts of the glowing road concept are drawn slightly blurred in a light green color, similar to HUD designs (see Fig. 7). Still images from the presented movies can be seen from in Fig. 4 to 7.

4.2 Discussion

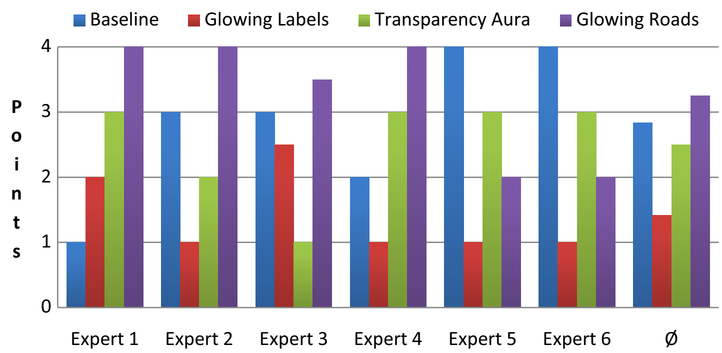

(a) Bird's eye.

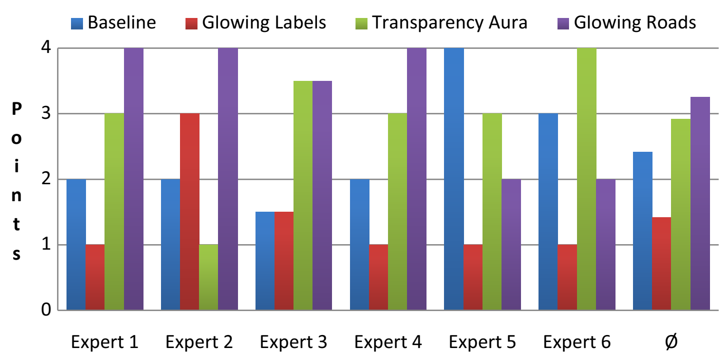

(b) Snail's view.

Fig. 8 Ranking of our concepts according to our six experts. Each concept was presented as a short movie. The concept glowing roads ranks first in both viewing perspectives.

In both views, glowing streets was ranked highest. 4 of 6 experts chose this as the best approach in both perspectives (bird and snail). Two experts stated that this concept improves orientation. Another expert liked how the glowing roads improve readability by creating an enhanced contrast to the background. One expert criticized the chosen color and suggested to continue the road in its original color. Finally, the last expert described this approach as being too colorful.

The second place is shared between the concept transparency label aura and our baseline. The former performs well in the snail's view, where labels are frequently hidden by 3D buildings. Our baseline sufficed in bird's eye view where occlusion plays a minor role and the spatial relationship is not needed.

Generally, the concept glowing labels was not approved and always ranked last. Three experts stated that the label seemed lost in the world and the coloring makes the visual association even more difficult. Two different experts did not approve that occluded parts should be marked with a different color. Finally, two experts criticized the color as being too vivid and distracting.

Our last concept, cutaways, was quickly dismissed by all experts, because it introduces too much animation. Every movement leads to new cut outs in the 3D buildings, thus removing parts of the world. When a lot of labels are present, the 3D world falls more and more apart.

When deciding which labeling space was best, 5 of 6 experts voted SS in bird's eye and 5 of 6 experts voted WS in snail's view. All but one expert agreed that in snail's view WS labeling was better despite the restricted readability.

4.3 Results

Concepts. As a first consequence, we dismiss two approaches: glowing labels and cutaways. In the experts' opinion, the disadvantages of the glowing labels concept (e.g. unaesthetical, bad visual association) outweigh the readability improvements. Cutaways introduce too much movement and destroy huge parts of the 3D world.
Visual association. Displaying the referenced feature besides the label is an important requirement for our implementation. One expert liked the transparency aura mainly because he was seeing the referenced road. The glowing labels ranked last because the association to the road becomes lost. In contrast, the concept glowing road recreates this reference.
Labeling technique. The last conclusion we draw, is the need to combine both SS and WS labeling in a 3D navigation. We choose SS in bird's eye and WS in snail's view. In snail's view the WS labels fits into the world's 3D space. In the bird's eye we hover at higher altitudes in which the world flattens. Therein, the better readability of 2D SS labels outweigh the deteriorated spatial relationship.

5 Implementation

We implement the selected concepts in an existing research platform for the visualization of navigation data. In this framework, the central processing unit (CPU) helps loading and preparing data for rendering. To ease the CPU load, our approaches run on the graphics processing unit (GPU) using shader programs.

5.1 Transparency label aura

In this concept, occluding parts of the buildings are faded out.

Overview. Our implementation consists of four steps. First, every building occluding a label is drawn into an offscreen buffer. In the second step, the entire set of buildings are again rendered offscreen. However, this time, we discard all fragments located in front of the occluded label – similar to an inverse depth test. In the third step, we combine these buffers to create a transparent aura around the label. Finally, we composite the result into the existing 3D world.

Implementation. The first rendering pass is trivial: we create an offscreen buffer and render all occluding 3D buildings into it. The second pass performs our inverse depth test in a fragment shader on the GPU. For this step, we need a texture (buffer) containing the depth information of all labels. We approximate each label with an object-oriented bounding-box (OOBB). And, because our experts stated in section 4.3 that the referenced objects should be seen, we slightly enlarge the bounding-box of each label. Then, we render all OOBBs of every visible label into a depth-only offscreen buffer. Finally, all buildings are drawn. In the fragment shader we compare the incoming depth value (of our buildings) $z_{building}$ with the depth value of our OOBBs (our labels) z_{label}. If $z_{label} > z_{building}$ the building occludes the label and we can discard this fragment. For the third step, we create a smooth blending in the transparency aura by rendering the OOBBs with a gradient texture. Finally, using this fullscreen alpha mask, we composite the results of the prior steps and render it over the current scene.

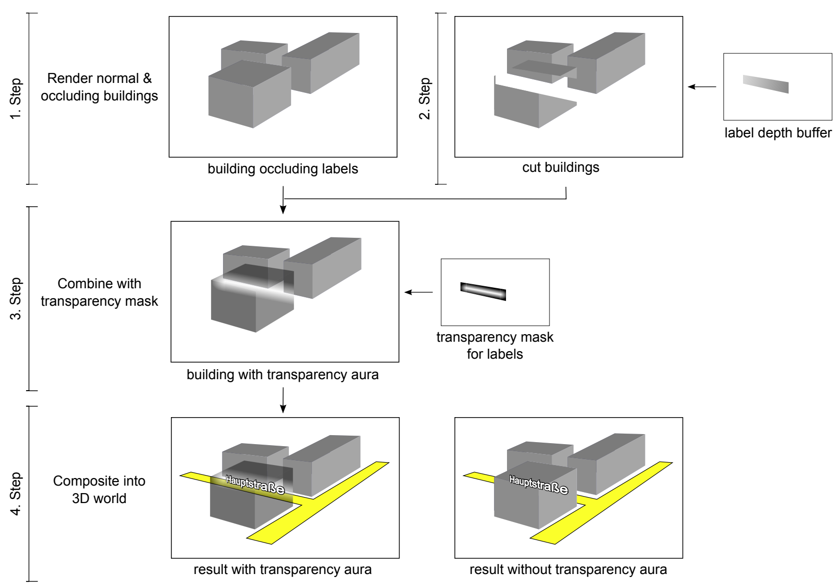

Fig. 9 GPU implementation of the transparency label aura approach.

5.2 Glowing streets

In this concept, all occluded parts of the roads are glowing over the 3D world.

Overview. The implementation consists of two steps. First, we detect which parts of the roads are being occluded. These parts are drawn with a selected color (e.g. light green). Then, optionally, a blurring filter is applied. Finally, the result is composited over the existing 3D world and all labels are rendered.

Implementation. Initially, we need the depth values of all rendered 3D buildings $z_{building}$ and roads z_{road}. Then, a fragment shader compares both depth values: If $z_{building} < z_{road}$, then the road is occluded and has to be drawn as a glowing road. If the glowing road is drawn with a single color, we simply output a constant color to an offscreen buffer. If we render the roads in their original color we first have to fetch this color. The resulting buffer can be smoothed with a blur shader and finally, composited with the existing 3D world. After these steps, all labels are drawn on top with a disabled depth test.

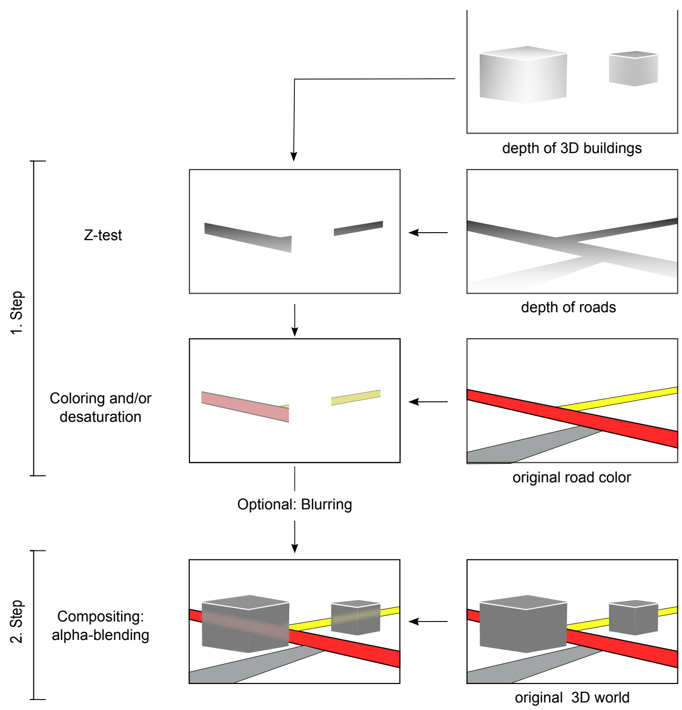

Fig. 10 GPU implementation of the glowing roads approach. Each step represents a shader.

6 Results

6.1 Benchmark

We benchmarked the approaches transparency aura and glowing roads. Our goal was to evaluate the performance scalability and suitability for real-world scenarios.
Configuration. The evaluation was done on an Intel Core 2 Duo E8400 3 Ghz CPU with 4GB RAM and Windows XP SP3. The GPU was a NVIDIA Quadro FX 580 (driver v275.89). To reduce the impact of data loading we preloaded all the needed data. Our performance measurement were done with a flight over a 3D city with roads, 3D buildings and labels. Fig. 12 shows the resulting performance graph during a flight of 20 seconds. We compare the baseline with the transparency aura and two variants of the glowing roads: using a single color and using the original road color. We measured the framerate for low 1024x768 (Fig. 12, top) and high resolution 1680x1050 (Fig. 12, bottom). During this run we tracked the number of buildings, road meshes and labels (see Fig. 12, middle).

Table 1 Average performance of the implemented concepts and framerate decrease (drop) compared to the baseline. Also, we list the performance impact when changing the resolution from 1024x768 to 1680x1050. We determine that both approaches are fillrate bound.

	framerate				
approach	1024x768	diff	1680x1050	diff	resolution impact
baseline	110 fps	–	59 fps	–	-46%
transparency label aura	82 fps	-25%	43 fps	-27%	-47%
glowing roads (single color)	90 fps	-18%	46 fps	-22%	-49%
glowing roads (road's color)	84 fps	-24%	42 fps	-29%	-50%

Results. At low resolution (1024x768) our new approaches behave similar to the baseline. Compared to our baseline, they incur a performance drop between 10-30%. The average performance decrease for every approach and for two resolutions can be seen in table 1. Our approaches are fillrate bound. At approximately twice the fragments (0.8 MP to 1.8 MP) we have a 50% performance decrease for every approach. Also, the increased number of 3D buildings, roads and labels do not impact the framerate as much as the increase in resolution (see Fig.12, middle).

6.2 User Study

Our goal was to evaluate the usability, attractiveness and novelty of our approaches.
Participants. We conducted an user study lasting 20 minutes with 24 persons aged between 17-45 consisting of 20 men and four women. About one third worked in the GIS domain. There were 9 students, 12 engineers, two programmers and one manager. Everyone had experience with commercial 3D navigation systems.

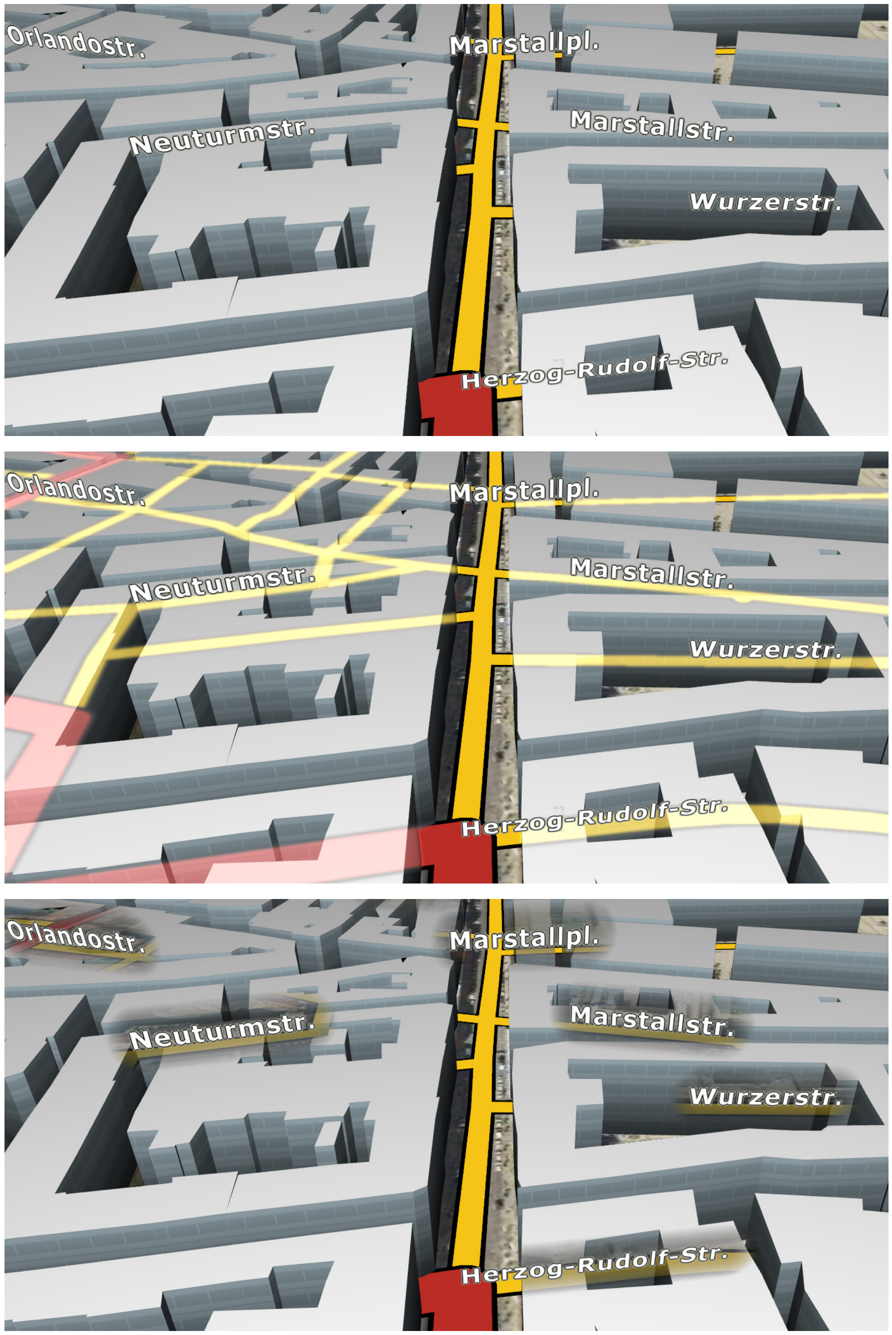

Fig. 11 Comparison of the implemented approaches in bird's eye with World-Space labeling: baseline (top), glowing roads (middle) and transparency label auras (bottom). As concluded from a conducted user study, the last two methods increase attractiveness and usability compared to the baseline.

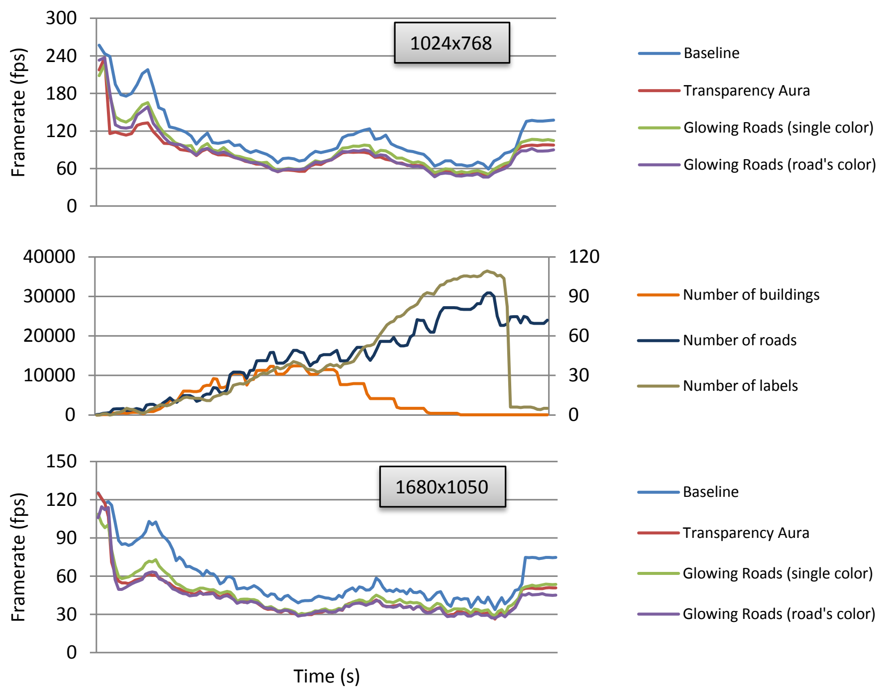

Fig. 12 Benchmark results of the GPU implementation: both approaches are fillrate-bound.

Study Design. These candidates tested the fully working prototypes of our baseline and the two implemented concepts: transparency label aura and glowing roads. In the first part of our evaluation, every subject flew three times the same 30 second lasting route through a 3D city. First, the baseline approach was active. Then, both new methods were shown in a changing order. After every flight the candidates had to fill out an AttrakDiff questionnaire (see Fig. 13). In the second part of the study, we let the subjects choose manually between all three concepts during a flight of two minutes. Finally, they completed a second informal questionnaire (see Fig. 14).

AttrakDiff. After experiencing the prototype, every candidate completed the AttrakDiff questionnaire from Hassenzahl et al. [14, 15]. They had to choose repeatedly between two different statements (e.g. attractive vs dull). These pairs were given by the AttrakDiff questionnaire to measure the perceived hedonic quality (HQ) and pragmatic quality (PQ). PQ is an indicator of the perceived usability of our concepts. HQ is divided into identity (HQ-I) and stimulation (HQ-S): HQ-I describes the user's identification, HQ-S defines the novelty of the tested concept. Finally, the questionnaire measures the overall attractiveness (ATT)

6.2.1 Results

Fig. 13(a) presents the averaged results of the AttrakDiff questionnaire. Compared to our baseline (orange), both approaches increase significantly every quality aspect and the overall attractiveness. The boxes in Fig. 13(b) indicate the overall classification in HQ and PQ. Therein, a placement in the top-right quadrant defines a very desired product. The size of the light boxes indicate the variability of the answers. In our case, the small box size of the baseline (orange) and glowing roads (blue) indicates a consistent opinion. In contrast, answers about the transparency aura (red) display more variation. In both figures, glowing roads (blue) achieve the best usability impact (PQ) and attractiveness (ATT). Overall, this validates the ranking of our experts from our pre-study.

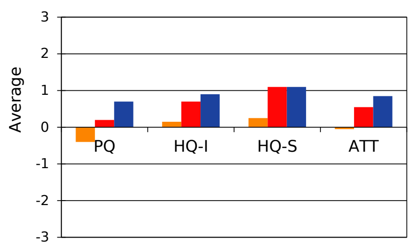

(a) Averaged values of the perceived qualities of the presented concepts.

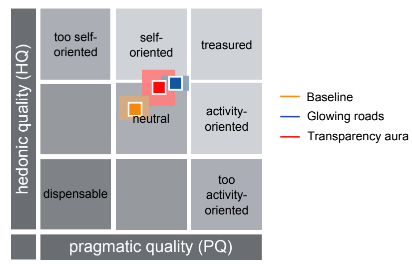

(b) Quality classification of the concepts and variability of the given answers.

Fig. 13 Resulting AttrakDiff questionnaire from our conducted user study. PQ describes the perceived pragmatic quality ($\approx$ usability), HQ-I the hedonic quality based on identity ($\approx$ user's identification), HQ-S the hedonic quality provided through stimulation ($\approx$ innovative) and ATT describes the concepts overall attractiveness. Compared to our baseline, both our presented approaches improve significantly the HQ, PQ and attractiveness.

Informal Questionnaire. Fig. 14 depicts the results of our second questionnaire. The majority state that the application of both approaches create an advantage compared to our baseline, create a better orientation and are aesthetically pleasing. The glowing roads display a higher distraction and are less calm than the transparency label aura. Our subjects would more likely use these approaches in a GIS than in a car. Overall, the proposed methods are perceived as a significant improvement compared to the baseline: 86% see transparency label aura and 77% glowing roads as enhancement.

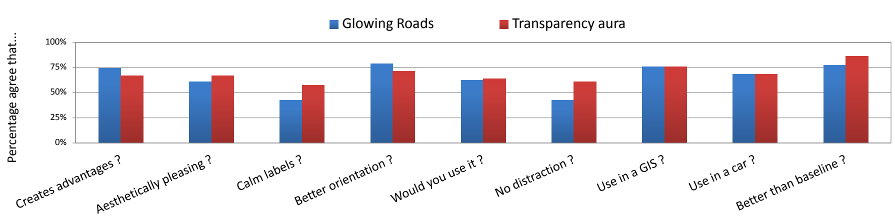

Fig. 14 Informal questionnaire answered by our 24 test candidates.

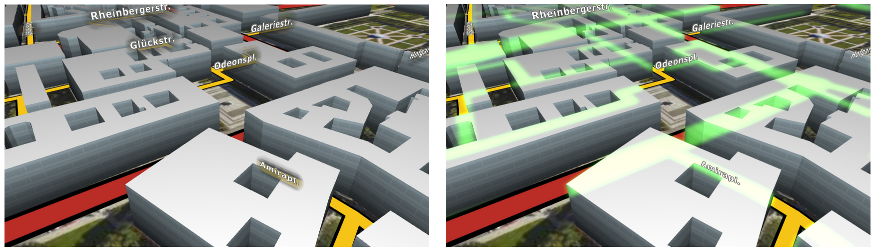

Fig. 15 Comparison of transparency label aura (left) and single colored glowing roads (right). Both figures are in bird's eye viewing space with WS labeling.

7 Conclusion

In this paper we have presented two new approaches, glowing roads and transparency label aura, which preserve the readability of occluded labels in 3D navigation maps while maintaining the reference to their corresponding object. We have described a prototypical implementation of both methods on the GPU running at interactive framerates. Our profiling has shown that these implementations are fillrate-bound. In a following user study including 24 subjects we compared them to our baseline: simply rendering all labels over the world, as done e.g. by Google Earth and almost every commercial navigation system. We have revealed that both our methods innovate and improve significantly the usability and overall attractiveness. Over 86% deem the approach glowing road better than our baseline. In further research, we plan to evaluate these approaches in real-world scenario, e.g. while driving through a city. Furthermore, a combination of both concepts could create new approaches, e.g. transparent road auras.

Acknowledgements We would like to thank Marco Matt for preparing and conducting the expert and user studies at our research labs. Also, we are grateful for the continuous feedback from Christopher Roelle while designing the labeling techniques. Finally, we thank Philipp Promesberger for developing the basis of the map rendering framework.

References

1. Ali, K., Hartmann, K., Strothotte, T.: Label layout for interactive 3d illustrations. Journal of the WSCG **13**(1), 1–8 (2005)
2. Bell, B., Feiner, S., Höllerer, T.: View management for virtual and augmented reality. In: Proceedings of the 14th annual ACM symposium on User interface software and technology, pp. 101–110. ACM (2001)

3. Bier, E.A., Stone, M.C., Pier, K., Buxton, W., Derose, T.D.: Toolglass and magic lenses: The see-through interface. In: Proceedings of the 20th annual conference on Computer graphics and interactive technique, pp. 73–80. ACM Press (1993)
4. Chen, J., Pyla, P., Bowman, D.: Testbed evaluation of navigation and text display techniques in an information-rich virtual environment. In: Virtual Reality, 2004. Proceedings. IEEE, pp. 181–289. IEEE (2004)
5. Cipriano, G., Gleicher, M.: Text scaffolds for effective surface labeling. Visualization and Computer Graphics, IEEE Transactions on **14**(6), 1675–1682 (2008)
6. Coffin, C., Höllerer, T.: Interactive perspective cut-away views for general 3d scenes. In: Proc. IEEE Symp. 3D User Interfaces (3DUI 06), IEEE CS, pp. 25–28. Press (2006)
7. Elmqvist, N., Assarsson, U., Tsigas, P.: Employing dynamic transparency for 3d occlusion management: Design issues and evaluation. In: Proceedings of INTERACT, ser. LNCS, C. Baranauskas, P. Palanque, J. Abascal, and S. D. J. Barbosa, Eds., vol. 4662, pp. 532–545. Springer (2007)
8. Fekete, J., Plaisant, C.: Excentric labeling: Dynamic neighborhood labeling for data visualization. In: Proceedings of the SIGCHI conference on Human factors in computing systems: the CHI is the limit, pp. 512–519. ACM (1999)
9. Goldstein, E.: Sensation and perception. Wadsworth Pub Co (2009)
10. Google Inc.: Google Earth v6.1.0.5001. (Software) (2012). URL http://earth.google.com/
11. Götzelmann, T., Ali, K., Hartmann, K., Strothotte, T.: Adaptive labeling for illustrations. Proceedings of Pacific Graphics 2005 pp. 64–66 (2005)
12. Götzelmann, T., Hartmann, K., Strothotte, T.: Agent-based annotation of interactive 3d visualizations. In: Smart Graphics, pp. 24–35. Springer (2006)
13. Hartmann, K., Ali, K., Strothotte, T.: Floating labels: Applying dynamic potential fields for label layout. In: Smart Graphics, pp. 101–113. Springer (2004)
14. Hassenzahl, M.: Attrakdiff (tm). URL http://www.attrakdiff.de/en/AttrakDiff/ (2007)
15. Hassenzahl, M., Burmester, M., Koller, F.: Attrakdiff: Ein fragebogen zur messung wahrgenommener hedonischer und pragmatischer qualität. In: Mensch & Computer, pp. 187–196 (2003)
16. Imhof, E.: Positioning names on maps. Cartography and Geographic Information Science **2**(2), 128–144 (1975)
17. Kalkofen, D., Mendez, E., Schmalstieg, D.: Interactive focus and context visualization for augmented reality. In: Proceedings of the 2007 6th IEEE and ACM International Symposium on Mixed and Augmented Reality, pp. 1–10. IEEE Computer Society (2007)
18. Kalkofen, D., Mendez, E., Schmalstieg, D.: Comprehensible visualization for augmented reality. Visualization and Computer Graphics, IEEE Transactions on **15**(2), 193–204 (2009)
19. Krüger, J., Schneider, J., Westermann, R.: Clearview: An interactive context preserving hotspot visualization technique. Visualization and Computer Graphics, IEEE Transactions on **12**(5), 941–948 (2006)
20. Maass, S., Döllner, J.: Dynamic annotation of interactive environments using object-integrated billboards. In: 14-th International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision, WSCG, pp. 327–334 (2006)
21. Maass, S., Döllner, J.: Efficient view management for dynamic annotation placement in virtual landscapes. In: Smart Graphics, pp. 1–12. Springer (2006)
22. Maass, S., Döllner, J.: Embedded labels for line features in interactive 3d virtual environments. In: Proceedings of the 5th international conference on Computer graphics, virtual reality, visualisation and interaction in Africa, pp. 53–59. ACM (2007)
23. Maass, S., Döllner, J.: Seamless integration of labels into interactive virtual 3d environments using parameterized hulls. In: 4th International Symposium on Computational Aesthetics in Graphics, Lisbon, pp. 33–40 (2008)
24. Maass, S., Jobst, M., Döllner, J.: Depth cue of occlusion information as criterion for the quality of annotation placement in perspective views. The European Information Society pp. 473–486 (2007)

25. Polys, N., Kim, S., Bowman, D.: Effects of information layout, screen size, and field of view on user performance in information-rich virtual environments. In: Proceedings of the ACM symposium on Virtual reality software and technology, November, pp. 07–09. ACM (2005)
26. Ropinski, T., Praßni, J., Roters, J., Hinrichs, K.: Internal labels as shape cues for medical illustration. In: Proceedings of the 12th International Fall Workshop on Vision, Modeling, and Visualization (VMV07), pp. 203–212 (2007)
27. Stein, T., Décoret, X.: Dynamic label placement for improved interactive exploration. In: Proceedings of the 6th international symposium on Non-photorealistic animation and rendering, pp. 15–21. ACM (2008)
28. Sygic: Sygic GPS Navigation. (Software) (2012). URL http://www.sygic.com/
29. Thielmann, T.: "‘you have reached your destination!"' position, positioning and superpositioning of space through car navigation systems. Navigation **2**, 27–62 (2006)
30. Viega, J., Conway, M.J., Williams, G., Pausch, R.: 3d magic lenses. In: ACM Symp. on User Interface Software and Technology, pp. 51–58 (1996)
31. Wolff, A., Strijk, T.: The map-labeling bibliography. Electronic Bibliography (2012). URL http://i11www.iti.uni-karlsruhe.de/ awolff/map-labeling/bibliography/