Investigation of Secondary Views in a Multimodal VR Environment: 3D Lenses, Windows, and Mirrors

Phanidhar Bezawada Raghupathy and Christoph W. Borst

University of Louisiana at Lafayette

Abstract. We investigate secondary view techniques in a multimodal VR environment for dataset exploration and interpretation. Secondary views, such as 3D lenses or mirrors, can present alternative viewpoints, different filtering options, or different data sets. We focus on 3D views showing surface features that are hidden in a main view. We present different view techniques, including new variations, and experimentally compare them. Experiment subjects marked paths on a geological dataset in a manner that required a secondary view for a portion of each path. We compared passive to interactive (reach-in) views, rotated to mirrored presentations, and box vs. window shapes. We also considered two types of path complexity arising from surface contact geometry impacting force feedback, as the level of lateral guidance provided by the contact geometry may impact relative effectiveness of different view techniques. We show several differences in task times, error rates, and subjective preferences. Best results were obtained with an interactive box shape.

1 Introduction

We investigate different ways of presenting a secondary view in a multimodal (multi-sensory) 3D environment and compare them for a path tracing task. We focus on secondary views displaying dataset regions hidden from the main view. In contrast to viewpoint or rendering changes in a main view, secondary views allow users to maintain a preferred view configuration and to simultaneously manage multiple projections, associated datasets, or filtering options (analogously to 2D windows). Understanding tradeoffs between different techniques and parameters will benefit VR-based scientific exploration applications such as geological interpretation.

Although various researchers considered secondary views for VR (summarized in Section 2), there has been little evaluation of their effectiveness. Even when present, such evaluations have not directly compared the various view techniques in VR. The main contributions of this paper are:

- We describe 3D secondary views, including variations not previously considered (e.g., a “reach-in” mirror, as opposed to a view-only mirror).
- We experimentally compare different secondary views and show:
  - Reaching in is very important: secondary views should include 3D interaction with viewed objects, not merely provide visuals.
  - Users prefer 3D boxes to window-like view shapes.

This is an author-formatted version. The original publication is available at www.springerlink.com. ISVC 2011, Part II, LNCS 6939, pp. 180-189.

- For marking areas hidden from the main view: there can be differences between mirrored and rotated view presentations, depending on other factors such as task motion direction and hand orientation.
- In a multimodal interface, there is interaction between view effects and task difficulty related to contact geometry and force feedback.

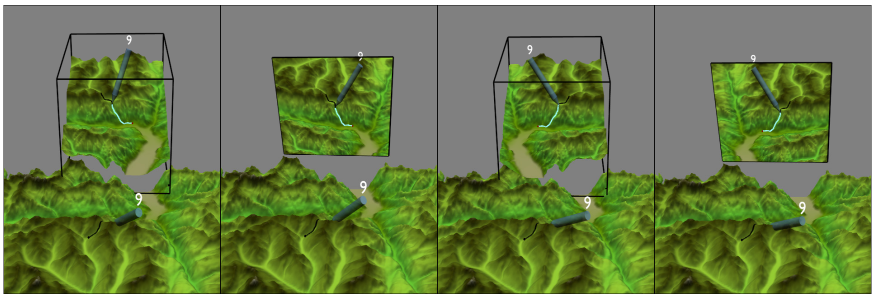

Fig. 1. Different secondary views. From left: mirrored 3D box view, mirrored window view, rotated 3D box view, and rotated window view

2 Related Works

Various researchers considered secondary views, calling them windows [2, 3, 4, 5, 8], boxes [6], lenses [7], or mirrors [9, 10, 11, 12,13]. Viega et al. [7] extended 2D lenses into 3D "Volumetric Lenses", where presentation in a box differed from surrounding view. Fuhrmann and Groller [6] refer to a similar concept as "Magic Box". Borst et al. [8] describe it more generally as a "Volumetric Window" for managing multiple viewpoints. For simplicity, we call these views 3D boxes (Fig. 1). We focus on views that show hidden sides of objects by presenting rotated or mirrored views.

Grosjean and Coquillart presented a "Magic Mirror" [9] analogous to a real mirror. Eisert et al. [10] and Pardy et al. [11] used virtual mirrors for augmented reality. König et al. [12] presented magic mirrors for volume visualization. Bichlmeier et al. [13] described a virtual mirror to reflect only virtual objects in augmented world. In this paper, we call these mirrored window views. We introduce "reach-in" mirrors and include an alternative called rotated window view (Fig. 1).

Some techniques like World-in-miniature [1], tunnel window [4] and SEAMS [2] provide interaction or reaching in for manipulating distant objects or navigating between different virtual worlds. In our work, we use reach-in to surfaces that were already reachable without secondary views but that can't be seen in the main view.

Elmqvist and Tsigas [14] classified many techniques for 3D occlusion management, including techniques affecting the main view. For example, Flasar and Sochor [15] compared navigation (active) techniques for manipulating objects behind obstacles. In our work, we focus on more passive techniques that avoid affecting the main view, with results more aimed at understanding 3D windowing approaches.

Numerous studies show that force feedback can affect performance. Typically these compare force feedback to no feedback, force-only to visual-only feedback etc., or haptic constraints to no haptic constraints. Some researchers, such as Faeth et al. [16], have used force feedback to aid operations on geological terrains. In our work, differences in force feedback are considered as they arise from different contact surface geometries and as effects may interact with view type.

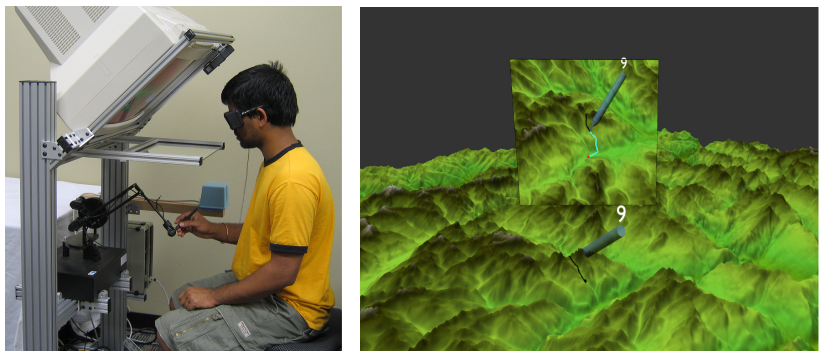

Fig. 2. Left: VR system with force stylus and mirror-based 3D display, a common setup for co-located visuals and force feedback. Right: User's view of terrain dataset and secondary view

3 Implementation Notes

Our multimodal environment (Fig. 2) renders visual and force feedback. Secondary views in 3D box shapes are rendered using techniques described in [8]. Secondary views with a window shape are rendered using standard stencil buffer mirror techniques, instead of texture mapping [9], to preserve depth and support reaching in.

Both box and window shapes auto-orient based on head position and a point of interest (POI). The POI can depend on context. For example, it may be the position of a pointer so that the view follows the pointer. In our experiments, we define a fixed POI as the center of a bounding box of a path being traced. This keeps the path centered and visible in the secondary view. 3D box and window views differ in the way content is seen and changes. 3D box content depends only on the POI and centers it in the box with constant box-relative orientation, related to traditional volumetric view rendering. But, for the window shape, different content can be seen depending on pose of the window and the POI, related to the usual way of rendering mirrors (although auto-orientation ensures that experiment paths are always centered in that view as well).

For a 3D box, a rotated view involves rotating box content 180 degrees around a local box-centered and box-aligned vertical axis and a mirrored view reflects the rotated view on a local horizontal axis. For window shapes, a rotated view rotates the original scene 180 degrees around a window-centered-and-aligned vertical axis and a mirrored view is obtained by reflecting the original scene about the window plane.

We automated view position to address manual placement bias for the experiment. Placement involves constraints with respect to a fixed reference coordinate system. Considering a fixed right-handed frame with X-axis being the VR display's rightward axis and Z-axis being its forward-facing axis (towards the user), a reasonable position (may not be optimal) for 3D box center (x, y, z - (depth of 3D box)/2) and window center (x, y, z) can be calculated as follows:

- x = x coordinate of center of a bounding volume of the path.
- y = highest y coordinate of point on the surface with respect to reference coordinate system plus (height of secondary view)/2 plus small positive offset.
- z = smallest z coordinate (farthest from user head) of the path minus the depth (z size) of a bounding volume of the path.

The Z offset makes the reach distance of window and 3D box approximately equal. The small Y offset moves the secondary view above the terrain.

For force rendering, we use a simple penalty-based method: force magnitude is proportional to stylus tip distance below a mesh surface, and force direction is the interpolated surface normal at the surface point directly above the stylus tip.

4 Experiment Methods

We conducted a within-subjects study comparing secondary views based on task time and error count (dependent variables) for a path tracing task. We also included a subjective preference session in which users compared certain conditions, switching a variable freely and indicating preference. For the objective portion, the independent variables, which produce 16 level combinations (conditions), are:

1. Reach mode (Reach-in, No reach-in)
2. Transform (Rotated, Mirrored)
3. Shape (3D box, Window)
4. Geometric guidance (With, Without)

When in Reach-in mode, a user reaches interactively into the secondary view to trace a path section, but otherwise the secondary view is used just for visual reference. Different levels of transform and shape were discussed in section 3. Path tracing may be supported by surrounding geometry (with geometric guidance) or not (without geometric guidance). For example, tracing along a valley or crevice results in lateral force-feedback cues that may help keep the stylus positioned along the path, while tracing along a flat portion or over ridges lacks this guidance.

Path tracing task: The task is representative of annotative marking for dataset interpretation. Although interpreters typically mark along features not yet marked, subjects traced an existing marked path to reduce cognitive and domain-specific aspects. Part of the each path was visible in the main view, but the remaining part was visible only in the secondary view. We were interested primarily in performance for tracing the hidden portion, including any time taken to transition between views. The study used four paths, shown in Fig. 3. Of the four paths, two have geometric guiding features in their hidden portions (left two of Fig 3) and the others do not. All paths had the same hidden portion length at the scale presented to subjects.

Hypotheses: Based on prior experience, we hypothesized that each independent variable was important and would impact performance or subjective preferences. We expected that reaching in improves speed and accuracy due to more direct interaction, even though transitioning between views involves extra time. We expected the 3D box view would be preferred for visual appearance but did not know if this would be reflected in performance. We speculated that mirrored view performs better than rotated view due to user familiarity with real-world mirrors. We expected that geometries producing lateral force guidance were easier to trace. Finally, we expected interactions, i.e., more notable effects when the task was difficult with respect to certain variables. For example, guidance would be more important when not reaching in.

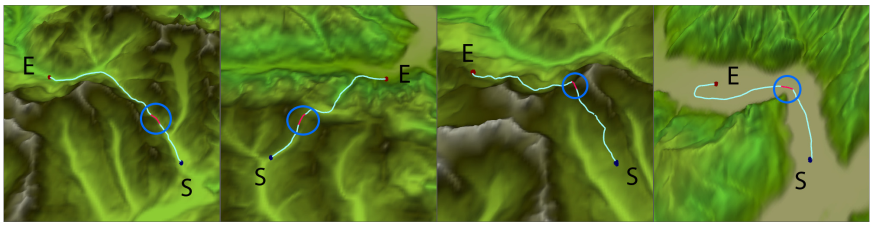

Fig. 3. Paths in the study (viewpoint differs from experiment viewpoint for clarity). “S” and “E” were added to the figure to show start and end points. Circles were added to show a transition point beyond which the path was no longer visible in the main view from subject’s perspective. All hidden portions have the same length

4.1 Apparatus

We used a mirror-based “fish tank” display as shown in Fig. 2 to co-locate haptic and visual feedback. Its horizontal mirror reflects a monitor image so users move a stylus directly in a virtual space below the mirror. Monitor resolution was 1024 x 768 with 100 Hz refresh rate divided into left/right frames by CrystalEyes glasses. Head position was tracked with an Ascension Minibird tracker synchronized to the monitor refresh to reduce jitter. A Sensable Phantom Premium 1.5 provided stylus input and force feedback. The host machine was a standard Dell graphics workstation.

4.2 Subjects

24 subjects participated. 19 were male and 5 were female. Ages ranged from 21 to 38, with an average of 26. 22 subjects were right-handed and 2 were left-handed. 8 subjects reported previous exposure to VR, 11 reported moderate to high experience with video games and 5 reported minimal or no experience with video games. Most subjects were students from computer science and engineering programs.

4.3 Main Experiment

The main experiment consisted of three sessions:

1. Practice (8 practice trials)
2. Session 1 (4 practice trials, 16 experimental trials)
3. Session 2 (4 practice trials, 16 experimental trials)

Experiment duration (including a subjective preference session) was typically 35-40 minutes. After Session 1, subjects were given a two minute pause.

4.3.1 Procedure

Per trial, subjects traced a path (Fig. 4) starting from a blue dot initially indicated by a blue arrow. The arrow vanished once the blue dot was contacted. The subject then traced the path and the contacted portion of the path turned black as it was traced. When the subject reached a pink-colored mark on the path, they switched focus to the secondary view, as the pink mark denoted the point after which the path was only visible in the secondary view. At that point, the subject either reached into the secondary view (reach-in) or used it only as visual reference (no reach-in) as directed by an arrow. The arrow disappeared once the subject reached in (reach-in) or passed the pink color (no reach-in). The subject then traced the remainder of the path to the end, denoted by a red dot, completing the trial. Additionally, whenever the subject moved off the path, recoloring stopped until the subject returned to the point where they left the path (threshold for both is 1.5mm from point on the path). Thus, there was no way to trace the path without moving through every point along it. Subjects were told to trace the path "as quickly as is comfortable". A counter at the end of the virtual stylus indicated elapsed time.

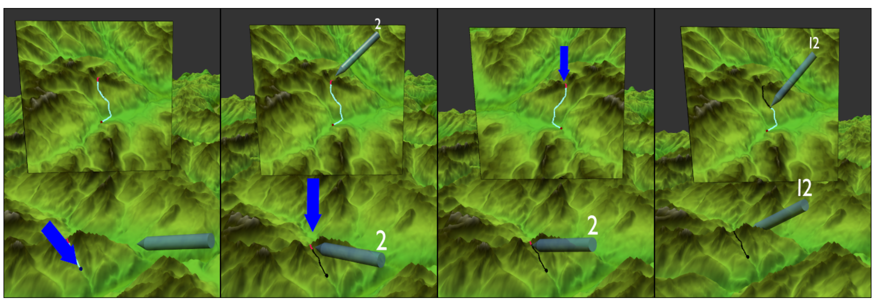

Fig. 4. Different stages of path tracing. From left to right: before starting, at transition for no reach-in, at transition for reach-in, and after transition for no reach-in

4.3.2 Condition Order and Randomization

The order of 16 conditions in each of Sessions 1 and 2 was randomized with the following constraints. We minimized switching of reach mode by requiring the first 8 trials of a session to be either all reach-in or all no-reach-in cases (random per subject). Two practice trials reminded subjects of reach mode after each switch. Within each resulting reach-in and no-reach-in set, the first 4 trials were all either rotated or mirrored cases (random per reach mode block). Within each rotated and mirrored set, the first 2 trials were all either 3D box views or window views (random per transform block). Within each resulting box or window view, there was one path with geometric guidance and one without (random per shape block). The two remaining paths appeared in the corresponding conditions during the other session.

4.4 Subjective Preference Experiment

In each of five preference trials following the main experiment, subjects indicated a preference after tracing a path and switching between techniques. Specifically, subjects compared reach-in to no reach-in cases (under randomized transform and shape), rotated reach-in to mirrored reach-in (shape randomized), 3D box reach-in to window reach-in (transform randomized), and no-reach-in versions of the latter two. Trial order was randomized per subject. Subjects could repeatedly trace the path and freely switch between relevant techniques by clicking a stylus button. Preference was indicated by a box click followed by a confirmation click.

5 Results and Discussion

5.1 Main Experiment Results

Task time was calculated as the amount of time taken to trace the hidden part of the path (including transition time). Error count was calculated as the number of times a subject moved off the path. Figures 5 and 6 summarize task time and error count means. We analyzed results with four-way repeated-measures ANOVA per dependent variable with Bonferroni correction for post-hoc tests.

Task time: Subjects traced paths faster when reaching in to secondary views than when using them only as a visual reference ($F(1, 23) = 51.002$, $p < .001$), with time averaging 56% shorter. Subjects were faster with 3D Box shapes than with window shapes ($F(1, 23) = 26.319$, $p < .001$), averaging 15% shorter task time. We detected no statistically significant effect of transform, overall, on task time $F(1, 23) = 0.202$, $p = .657$). Lateral geometric guidance improved task times ($F(1, 23) = 31.849$, $p < .001$), with average 34% task time reduction over other paths (same path lengths).

There were significant reachmode-guidance and transform-guidance interactions, ($F(1, 23) = 14.790$, $p < .001$) and ($F(1, 23) = 9.504$, $p = .005$), respectively. We investigated interactions with reduced-variable ANOVAs at fixed levels of variables of interest. The increase in task time from guidance to no guidance averaged 20% for reach-in compared to 68% for no-reach-in, indicating guidance was especially important in the more difficult case of no reach-in. For transform-guidance interaction, mean task time for mirrored view was shorter with geometric guidance and longer without geometric guidance when compared to rotated view.

Error Count: All independent variables affected error count. Subjects stayed on the path better when reaching in to a secondary view than when it was only a visual reference ($F(1, 23) = 113.26$, $p < .001$), with error count averaging 45% smaller. 3D Box shapes were better than window shapes ($F(1, 23) = 16.691$), with error count averaging 12% smaller. Overall, rotated views produced fewer errors than mirrored views ($F(1, 23) = 9.986$, $p = .004$), averaging 13% lower. Geometric guidance reduced errors ($F(1, 23) = 34.779$, $p < .001$) by an average of 25%.

There were significant reachmode-guidance and transform-guidance interactions, ($F(1, 23) = 18.049$, $p < .001$) and ($F(1, 23) = 20.456$, $p < .001$), respectively. We investigated these as we did for task time interactions. Again, guidance was more important in the no-reach case: there was a significant effect of geometric guidance for no-reach-in (F (1, 23) = 92.7, p < .001) but not for reach-in mode (F (1, 23) = .772, p = .389). And, transform was more important for no-guidance cases: there was a significant effect of transform for no-guidance cases (F (1, 23) = 26.095, p < .001), but not for guidance cases (F (1, 23) = 0.087, p = .771).

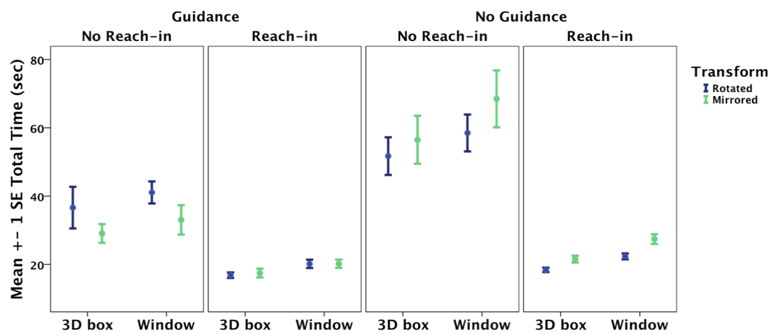

Fig. 5. Task time means and standard error bars for the 16 conditions

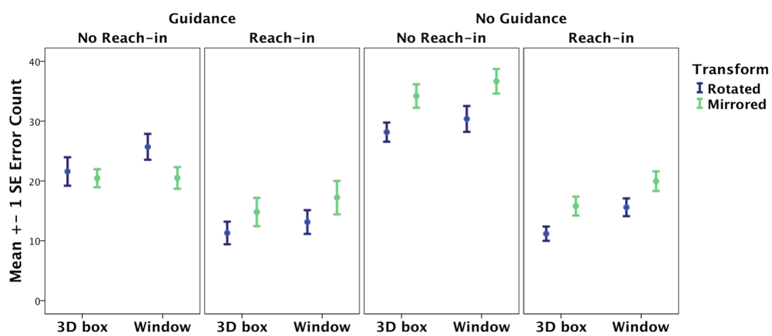

Fig. 6. Error count means and standard error bars for the 16 conditions

5.2 Subjective Preference Experiment Results

For each subjective preference question, each subject was given a score of zero or one depending on the technique selected. We used one-parameter two-tailed z-tests to detect significant difference in mean score from 50% (the no-preference score).

For reach mode: Significantly, all 24 subjects preferred reach-in to no reach-in.

For transform: Significantly, 17 Subjects preferred the mirrored to rotated transform in no-reach-in cases (Z (24) = 2.041, p = .041). For reach-in cases, there was no statistically significant preference, with 12 subjects preferring each technique.

For shape: Significantly, 20 Subjects preferred 3D box shapes to window shapes in reach-in (Z (24) = -3.326, p < .001) and 17 Subjects preferred 3D box shapes to window shapes in no reach-in cases (Z (24) = -2.041, p = .041).

5.3 Discussion

Our hypotheses are largely supported by results except that performance measures did not consistently favor mirror over rotated transform.

The most promising secondary view is a reach-in 3D box. Regarding transform type (mirrored vs. rotated), subjective preference and objective results differ, and there may be other factors to consider. For example, in some applications, interpreters may want to see a view that preserves "handedness" of data, which is violated by mirrored views but not by rotated views.

Even though reaching in to a secondary view requires additional transition time, both task time and error count were still reduced significantly by reaching in (averaging 56% and 45%, respectively), and subjective preference results unanimously supported reach-in mode over no-reach. There are two aspects that make the task more difficult without reaching in: the interaction is less direct (not co-located), and the secondary view that is being used is flipped along some axis with respect to the required hand motion (for both mirrored and rotated cases).

3D box shapes were better than window shapes in terms of task time, error count, and subjective preference. Note the 3D box technique provides more consistent content while window versions are more sensitive to position and orientation (Section 3). This makes placing windows more difficult, as there can be substantial deviation in viewed path orientation and depth with relatively small window position and orientation changes. Although we believe our window placements were good and well-matched to the 3D box versions, we cannot be sure they were optimal, and this illustrates the problem of sensitivity to placement. In real applications, there can be aspects preventing ideal placement, such as occluding objects or differences between ideal locations for reaching in (for comfortable depth) and ideal placement for viewed content orientation (so visuals match hand motion). Sensitivity to viewpoint is also a problem for head-tracked VR.

The lateral guidance from certain geometric features helps performance, but the extent depends on other view parameters. Paths with guidance significantly averaged 34% faster and 25% less error-prone. Geometric guidance had a stronger influence when subjects could not reach in to the secondary view (i.e., when the task was more difficult).

Subjectively, subjects preferred mirrored over rotated views in no-reach mode but showed no preference in reach-in mode. Objective performance measures contrast this by showing the rotated view had lower error count. We believe that factors not explicitly studied affect performance results. For example, the position and direction of paths and handedness of subjects (affecting pen tilt) may be influential.

6 Conclusions and Future Work

We discussed secondary views in a multimodal environment to overcome visual constraints (hidden features), and we compared different secondary views based on reach mode, transform, and shape. Our study confirmed that a 3D box view with reach-in interaction was the best considered secondary view for a hidden path tracing task, and a mirrored view appeals to users when not reaching in. Surface geometry impacts performance, particularly when users don't reach in: features that result in good lateral force cues help users overcome the indirect nature of no-reach interaction. For real terrain marking applications, the presence of such features hinges on the specific interpretation task, so it is important to optimize other view parameters for tasks where these features are lacking.

Even though a mirrored secondary view was preferred based on subjective comparisons, the performance of mirrored and rotated views should be further studied by careful consideration of path orientations and right- and left-handed subjects.

References

1. Stoakley, R., Conway, M., Pausch, R.F.: Virtual Reality on a WIM: Interactive Worlds in Miniature. In CHI (1995) 265-272
2. Schmalstieg, D., Schaufler, G.: Sewing Worlds Together with SEAMS: A Mechanism to Construct Complex Virtual Environments. Presence (1999) 449-461
3. Robertson, G.G., Dantzich, M.V., Robbins, D.C., Czerwinski, M., Hinckley, K., Risden, K., Thiel, D., Gorokhovsky, V.: The Task Gallery: a 3D window manager. In CHI (2000) 494-501
4. Kiyokawa, K., Takemura, H.: A Tunnel Window and Its Variations: Seamless Teleportation Techniques in a Virtual Environment. In HCI International (2005)
5. Ware C., Plumlee, M., Arsenault, R., Mayer, L.A., Smith, S.: GeoZui3D: Data Fusion for Interpreting Oceanographic Data. OCEANS (2001) 1960-1964
6. Fuhrmann, A.L., Gröller, E.: Real-time techniques for 3D flow visualization. In IEEE Visualization (1998) 305-312
7. Viega, J., Conway, M., Williams, G.H., Pausch, R.F.: 3D Magic Lenses. In ACM Symposium on User Interface Software and Technology (1996) 51-58
8. Borst, C.W., Baiyya, V.B., Best, C.M., Kinsland, G.L.: Volumetric Windows: Application to Interpretation of Scientific Data, Shader-Based Rendering Method, and Performance Evaluation. In CGVR (2007) 72-80
9. Grosjean, J., Coquillart, S.: The Magic Mirror: A Metaphor for Assisting the Exploration of Virtual Worlds. In SCCG (1999) 125-129
10. Eisert, P., Rurainsky, J., Fechteler, P.: Virtual Mirror: Real-Time Tracking of Shoes in Augmented Reality Environments. In ICIP (2) (2007) 557-560
11. Pardhy, S., Shankwitz, C., Donath, M.: A virtual mirror for assisting drivers. In IV (2000) 255-260
12. König, A., Doleisch, H., Gröller, E.: Multiple Views and Magic Mirrors - fMRI Visualization of the Human Brain. In SCCG (1999) 130-139
13. Bichlmeier, C., Heining, S.M., Feuerstein, M., Navab, N.: The Virtual Mirror: A New Interaction Paradigm for Augmented Reality Environments. IEEE Trans. Med. Imaging (2009) 1498-1510
14. Elmqvist, N., Tsigas, P.: A Taxonomy of 3D Occlusion Management Techniques. In VR (2007) 51-58
15. Flasar, J., Sochor, J.: Manipulating Objects Behind Obstacles. In HCI (14) (2007) 32-41
16. Faeth, A., Oren, M., Harding, C.: Combining 3-D geovisualization with force feedback driven user interaction. In GIS (2008)