Egocentric Space-Distorting Visualizations for Rapid Environment Exploration in Mobile Mixed Reality

[1]Christian Sandor* [1]Andrew Cunningham† [1]Ulrich Eck‡ [1]Donald Urquhart§ [1]Graeme Jarvis¶
[1]Arindam Dey‖ [1]Sebastien Barbier** [1]Michael R. Marner†† [2]Sang Rhee‡‡

[1]University of South Australia
Wearable Computer Laboratory
Magic Vision Laboratory

[2]Kyungnam University
Department of Computer Science and Engineering

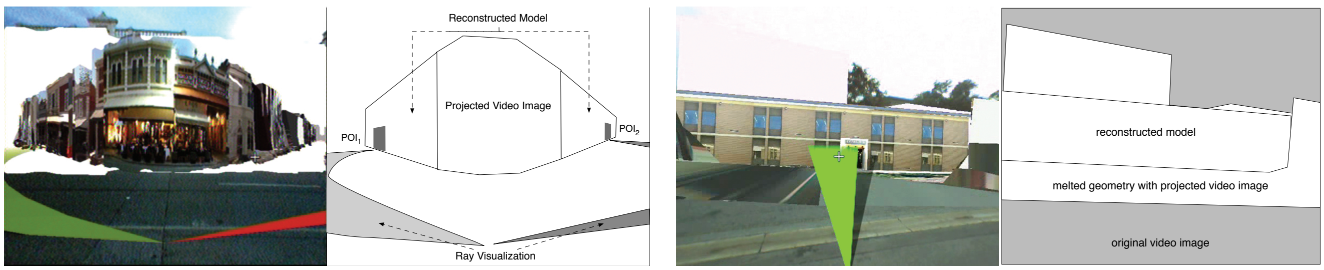

(a) (b) (c)

Figure 1: Our space distortion visualizations enable users to rapidly grasp the location of points of interest (POIs) that are either outside their field of view or otherwise occluded. (a) illustrates Radial Distort, revealing POIs that are outside the field of view of the user. (b) illustrates Melting, for discovering occluded POIs.

Abstract

Most of today's mobile internet devices contain facilities to display maps of the user's surroundings with points of interest embedded into the map. Other researchers have already explored complementary, egocentric visualizations of these points of interest using mobile mixed reality. Being able to perceive the point of interest in detail within the user's current context is desirable, however, it is challenging to display off-screen or occluded points of interest.

We have designed and implemented space-distorting visualizations to address these situations. While this class of visualizations has been extensively studied in information visualization, we are not aware of any attempts to apply them to augmented or mixed reality. Based on the informal user feedback that we have gathered, we have performed several iterations on our visualizations. We hope that our initial results can inspire other researchers to also investigate space-distorting visualizations for mixed and augmented reality.

Index Terms: H.5.1. [Information Interfaces and Presentation]: Multimedia Information Systems—[Artificial, augmented and virtual realities] I.3.6 [Computer Graphics]: Methodology and Techniques—[Interaction Techniques]

*christian@sandor.com
†andrew.cunningham@unisa.edu.au
‡ueck@net-labs.de
§swiftless@gmail.com
¶fanged@gmail.com
‖Arindam.Dey@postgrads.unisa.edu.au
**sebastien.b@rbier.com
††marnermr@cs.unisa.edu.au
‡‡syrhee@kyungnam.ac.kr

1 Introduction

Most of today's mobile internet devices contain facilities to search the user's immediate environment for points of interest (POIs), with the search results displayed on a map. Several research projects have aimed to provide more intuitive visualizations of POIs by using augmented reality to overlay symbol on a video image of the real-world. In one recent example presented by Nokia [9], users can use their mobile phone as a magic lens on the environment. This egocentric information display provides perceptual advantages that allow users to grasp spatial relations faster than an exocentric map display [11].

Egocentric visualizations, however, introduce several new challenges; how to display POIs that are, 1) behind the user, 2) outside of the user's field of view (FOV), or 3) occluded by real-world objects? The core idea of our research is to apply *space-distorting visualizations* to these situations. While this class of visualizations has been extensively studied [4] in information visualization, we are not aware of any attempts to apply them to augmented or mixed reality (MR). To show points of interest outside of the user's FOV, we employ a technique we call *Radial Distort* (Figure 1a). To show occluded points of interest, we *Melt* the occluding objects (Figure 1b). Finally, to make our visualizations more understandable for users, we provide cues about distortions through bending *Rays* and live *Projected Video* on deforming geometry.

Contribution The core contribution of this paper is the initial exploration of space-distorting visualizations for mobile MR. We have designed and implemented two complementary visualization techniques (Melt, Radial Distort) and two supporting cues (Rays, Projected Video). Projected Video is a novel cue that is specifically useful for space-distorting visualizations in mobile MR, as real-world context is maintained as much as possible while distorting. Based on user feedback, we have extended our visualizations in two

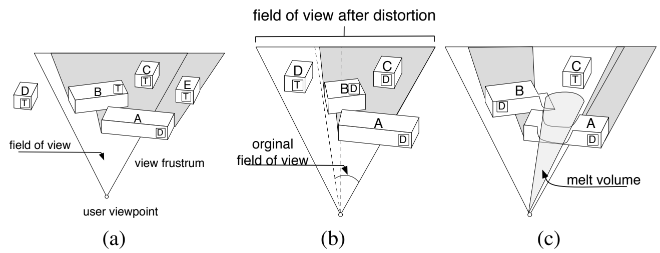

Figure 2: Our adaption of the Elmqvist and Tsigas occlusion model and how they apply to Melt and Radial Distort. (a) Elmqvist and Tsigas occlusion model to include targets outside the FOV. Target objects are flagged with “T” and distractors are flagged with “D”. (b) Shows a schematic of the Radial Distort after distortion. (c) The melt volume with distractor objects flattened.

ways: First, the Ray cue, which illustrates the state of the distortion, was added. Second, zooming after melting was added in order to make the revealed POIs more legible. A side contribution of this paper is the extension of the taxonomy of occlusion management techniques presented by Elmqvist and Tsigas [5] to describe POIs that are outside the user’s FOV.

The rest of this paper is structured as follows: First, we discuss related work. We then describe our visualization techniques and present their implementation. We conclude by presenting initial user feedback and future work.

2 Related Work

MR can enhance user’s perception of their environment by showing information they cannot directly sense [6], however designing such techniques is not trivial. We recognize the need for MR techniques that display detailed POIs that are either occluded or outside their FOV within the user’s current context. Space distortion is an accepted information visualization technique for providing local *detail* and global *context*; a focus+context technique. Space distortion visualizations leverage our visual gestalt abilities to present multiple POIs continuously in their global context, thus reducing cognitive load. Carpendale et al. [4] extend 2D space distortion schemes into 3D space, noting that such techniques can provide a clear path to the target, maintain aspects of the surrounding context, and scale to multiple foci. Vallance [12] applies space distortion to 3D planar navigation, where a multi-projection renders user’s immediate context as an egocentric view, and distant information as an exocentric view. The authors argue that such an approach benefits by removing the cognitive context switch that map navigation requires.

Bell et al. [2] use popup annotations to point to on-screen and off-screen objects. A worlds in miniature provides a level of context for the user’s view within the environment. These cues indicate both distance and direction, but the context and focus are not continuously integrated, which implies this is an overview + detail technique [3]. Disparate overview and detail displays can increase cognitive load due to the cost of visual search and working memory when switching contexts. Güven and Feiner [8] describe a lifting metaphor for viewing occluded POIs. Users are provided cues, such as replicas of real-world buildings, to help maintain context. However, lifting provides less context than our techniques, since POIs are removed from their immediate real-world context and viewed from below, making recall difficult.

We apply space distortion visualization techniques to display POIs that are outside the FOV or occluded. Elmqvist and Tsigas describe the most thorough taxonomy for 3D occlusion management techniques [5]. Their taxonomy is described through a model of occlusion (Figure 2(a)), where an object is either a target, analogous to a point of interest, or a distractor, containing no intrinsic value. An object o is occluded if a line segment r from the user’s eye to the target passes through another object. In this model, the authors identify five occlusion management design patterns. Of immediate interest to our mobile MR research are *Virtual X-Ray*, *Volumetric Probe* and *Projection Distorter* since they are most applicable to an egocentric visualization, whereas *Multiple Viewports* and *Tour Planner* are exocentric.

Previously, we have described a Virtual X-Ray technique for visualizing occluded targets [1] in a mobile MR environment. This technique shows occluded points of interest rendered in their correct location, with highlighted edge cues from the video image aiding depth perception. To complement this depth perception, we suggest a cutout technique that cuts a hole through the virtual geometry. The Melt technique described in this paper can be considered a natural progression of the cutout technique. Both fit the volumetric probe design pattern, suited to retrieving visual information associated with a POI.

We extend the Elmqvist model (Figure 2(a)) to consider objects that are outside the user’s FOV (D) or occluded and outside the FOV (E). A line segment r originating at the viewpoint is considered outside the FOV if the angle made by the segment and the user’s viewing direction is greater than $1/2$ FOV or less than $-1/2$ FOV. An object is considered outside the FOV if all r to the object are outside the FOV. With this expanded model, our radial distort technique fits the projection distorter design pattern, suited towards POI discovery.

3 Visualizations

We explore space distortion for representing POIs that are either outside the user’s FOV, occluded, or both outside the FOV and occluded — situations that we feel have not currently received enough attention in MR based visualizations and are pertinent to the task of environment exploration. Space distortions present target POIs within the user’s current view, representing both the local detail of POIs, as well as their relationship to the global context (being the user’s point of view). Such attributes are ideal for environment exploration as the user is not required to move their point of view in order to build a strong spatial mental model of their environment and its relationship to the POIs.

Exploring 3D space distortions in mobile MR with any veracity requires a 3D reconstruction of the real-world to be distorted. To blur the line between this reconstruction and reality, we use a *Projected Video* cue to project the real-world video image from the user’s current point of view onto the reconstructed model. When the reconstructed model is distorted, the video image is distorted accordingly, minimizing the dissonance between the real-world and reconstructed model, and therefore reducing the cognitive load required to understand the distorted space.

To further reduce cognitive load, we use a *Ray* cue. The Rays act as an overview and unifying signpost for the visualizations — a “cognitive anchor” for space distortion. Rays are rendered as wedges, emanating from the user towards the point of interest. Rays that pass through objects such as buildings become semi-transparent as an added depth cue. When distortions such as the Radial Distort occur, the Rays bend towards the distortion, indicating both their original direction in the near half of the ray and the distorted location in the far half. The amount of distortion is reinforced in the ray color, akin to stress visualization, where a ray under no distortion is green and a heavily distorted ray is rendered red. These ray cues were added from informal user feedback and are invaluable in providing a grounding for users who have no experience with space distortions.

Our visualization techniques build upon the aforementioned cues to distort the space while keeping semblance with the real-world. To bring POIs outside the user’s FOV into view, we radially dis-

Figure 3: Video images of our visualization techniques used in a city location. We enable users to rapidly grasp the location of points of interest that are outside their FOV by using our Radial Distort visualization. (a-c) show three frames of the animated Radial Distort visualization. (d-f) Show that occluded locations can be observed by our Melt visualization: the occluding geometry is flattened while projecting the video onto it. (e-h) POIs are zoomed closer to the user for visibility.

tort the world around the user, compressing regions of the FOV that don't contain a POI. To reveal occluded POIs, we melt the occluding geometry. To reveal a POI that is both outside the FOV and occluded, we first radially distort the POI into view and then melt the occluding objects.

3.1 Radial Distort

The purpose of Radial Distort is to bring targets outside the FOV within the FOV, as illustrated in the Figure 2(b). Figure 3(a-c) show a video sequence of the visualization where one POI is outside the FOV. This POI is rotated inwards, so that it lies at the border of the FOV. The center of the FOV is compressed accordingly. The video image is projected in realtime onto the deforming geometry.

The benefit of such a technique is the ability to bring a POI to the attention of the user without forcing them to change their viewing direction, thereby keeping current context. Users can then turn their view towards the POI, with the distortion interactively unravelling as they do. In our experience, this interaction is more natural than using arrows to point to offscreen content.

Like Vallance [12], Radial Distort presents both immediate detail and continuous global context. Exact spatial relationship is degraded in favor of representing relative location to the user's current point of view. This follows the projection distorter design pattern. The user's current point of view is merged into n points of view for each point of interest, benefitting POI discovery at the cost of POI invariance.

3.2 Melt

The purpose of Melt is to reveal occluded targets. Occluded targets are revealed by virtually melting the distractor objects (Figure 3(d-f)). A circle sector volume originating from the user in the direction of the POI defines the melt volume (Figure 2(c)). The POI may be zoomed to gain more screen space and therefore present more information about its immediate context (Figure 3(g-h)). Melt fits the volumetric probe design pattern [5]; however, it is passive/offline. The user selects the target from a list and the system animates the melting of distractors within the melt volume.

The melting metaphor replicates a common metaphor which we reinforce through animation. Like Virtual X-Ray, which is inspired by the 'superman'-like ability to see through buildings, Melt is inspired by the superman-like ability to melt buildings. Compared to Virtual X-Ray, Melt is suited to situations where there is more than one occluding object, such as POIs that are two or more streets over from the user. In this case, Virtual X-Ray loses most depth information besides the most immediate distractor. Virtual X-Ray also introduces high visual complexity when rendering transparency, decreasing depth perception and increasing cognitive load.

4 IMPLEMENTATION

Implementing space distortion visualizations in a mobile MR environment is novel, and we faced two notable challenges in doing so. Firstly, a reconstructed model of the real-world is needed and secondly, distortions must be calculated and rendered in realtime. We implemented a pipeline for creating a reconstruction of the real-world and our visualizations are implemented as GLSL shaders (Section 4.1). This approach allowed us to alter our testing environment without requiring expensive remodeling, and allowed us to rapidly prototype many space distortions without expensive recompilation.

For our techniques to work, we require a textured 3D model of the world. Our modeling pipeline takes a point cloud of a city (in our case, we obtained a pointcloud of Adelaide's city center from the Airborne Research Australia at Flinders University, South Australia). Next, we apply Poullis' and You's algorithm [10] for creating a mesh from the pointcloud. We developed an algorithm to texture this mesh using spherical panoramas and their associated metadata from Google's Streetview [7]. Finally, we have manually added tags to our model to delineate POIs, however, in the future this too may be automated from the various geographical information systems available.

Our visualizations run on a belt-worn computer with a 7" handheld display. The computer features a Pentium-M 2.0 GHz, 2 GB RAM, and NVIDIA GeForce 6600 graphics chipset. The handheld display is a Xenarc 7 VGA touchscreen, to which we have attached an InterSense InertiaCube3 and a Fire-i camera. It is important to note that our system does not provide location tracking. The InertiaCube3 is used to provide rotation information, while the users position must be provided manually. Our software platform is implemented in the Python programming language. The visualizations, being implemented as a shader program, have a minimal impact on performance. We achieve approximately 20fps running our visualizations on a city model with more than 50,000 vertices.

4.1 Visualization Implementation

Visualizations are implemented as GLSL 1.2 shaders. The video projection cue, computed in the vertex shader, projects texture coordinates for the video image onto the front-most vertices of the undistorted scene. For each vertex in the world model, the homogenous coordinate depth is compared to the depth buffer of the undistorted scene. If the depths are equal (compensating for error), then the vertex receives a projected texture coordinate for the video texture.

The next stage distorts vertex positions to show melting and radial distort. In our implementation, we describe all POIs $p \in P$ as angles p_{angle_min} and p_{angle_max}, relative to the user's viewing direction, and a distance p_z, relative to the user's current position.

Algorithm 4.1: VERTEX SHADER(V, P, d)

for each $v \in V$
  do
    // calculate γ, the angle of v relative to the user's current
    // viewing direction d
    $\gamma \leftarrow \text{ANGLE}(v, d)$

    // radial distortion
    $\gamma \leftarrow \gamma * \text{RADIALDISTORT}(\gamma)$

    // melt any occluding objects
    for each $p \in P$
      do **if** $p_{angle\_min} \geq \gamma \leq p_{angle\_max}$ **and** $v_z \leq poi_z$
        then $v_z \leftarrow 0$

Where V is the set of all vertices and RADIALDISTORT returns an angular coefficient that radially distorts the vertex into view. Our RADIALDISTORT implementation linearly compresses the FOV until all POIs are visible, and is defined by the following equation:

$$r = \begin{cases} \frac{½FOV}{\arg\min(p_{angle\_min} | p \in P)} & \text{for } \gamma \leq 0, \\ \frac{½FOV}{\arg\max(p_{angle\_max} | p \in P)} & \text{for } \gamma \geq 0. \end{cases}$$

Where r is the angular coefficient returned by RADIALDISTORT, such that it is the ratio of half the FOV to the greatest angle of a POI outside the FOV. At this stage, all visualizations have been calculated and are rendered by the fragment shader.

5 Conclusions and Future Work

In this paper, we have presented our initial explorations of space-distorting visualizations for mobile MR. We extend the Elmqvist and Tsigas taxonomy of occlusion management to consider POIs that are outside the FOV and describe our techniques in that model. Our visualizations enable users to see POIs that are occluded and outside the FOV within their current real-world context.

After completing an initial prototype of our visualizations, we tested and refined them for ten days in the streets of the city of Adelaide, South Australia. We have shown our visualizations repeatedly to pedestrians that happened to be near our experiments. Based on their feedback, we added zooming to the melt visualization, as users had complained that the locations revealed were often difficult to see due to distance, and also introduced the ray cues to aid users who have no experience with space distortions.

This informal feedback has been invaluable in designing the visualizations, however, we consider a formal evaluation to be the most important piece of future work. The primary goal of these evaluations will be to collect quantitative information including navigation speed (how fast can users navigate from A to B using our system), recall percentage (how persistent are the memories of the information retrieved with our system), and error rate (how many misinterpretations by the user happen).

On the technical side, Our current prototype works only on fixed positions. For conducting our initial explorations, this was sufficient. However, in order to conduct formal user studies, we are currently working on an edge-based tracker in order to provide pixel-accurate augmentations.

Our visualizations could assist users in a variety of situations. However, there are several cases for which we haven't yet found the right space distortion technique; for example, POIs that face away from the user are not easy to display. A similar problem is that we can't yet visualize points of interest behind the user. Only when we have found appropriate visualizations for these situations, will we have completed our visualization toolbox.

Acknowledgements The authors wish to thank Jorg Hacker from Airborne Research Australia for providing a pointcloud of Adelaide; Charalambos Poullis and Suya You (University of Southern California) for meshing it. Last, but not least, we would like to thank our colleagues at the Wearable Computer Laboratory, who have helped with many parts of this project.

References

[1] B. Avery, C. Sandor, and B. H. Thomas. Improving spatial perception for augmented reality x-ray vision. In *Proceedings of the IEEE Virtual Reality Conference 2009*, Lafayette, Louisiana, USA, March 2009.

[2] B. Bell, T. Höllerer, and S. Feiner. An annotated situation-awareness aid for augmented reality. In *UIST '02: Proceedings of the 15th Annual ACM Symposium on User interface Software and Technology*, pages 213–216, Paris, France, 2002. ACM Press.

[3] S. K. Card, J. D. Mackinlay, and B. Shneiderman, editors. *Readings in information visualization: using vision to think*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1999.

[4] M. Carpendale, D. Cowperthwaite, and F. Fracchia. Extending distortion viewing from 2 D to 3 D. *IEEE Computer Graphics and Applications*, 17(4):42–51, 1997.

[5] N. Elmqvist and P. Tsigas. A taxonomy of 3d occlusion management for visualization. *IEEE Transactions on Visualization and Computer Graphics*, 14:1095–1109, 2008.

[6] C. Furmanski, R. Azuma, and M. Daily. Augmented-reality visualizations guided by cognition: Perceptual heuristics for combining visible and obscured information. In *Proceedings of the 1st International Symposium on Mixed and Augmented Reality*. IEEE Computer Society Washington, DC, USA, 2002.

[7] Google Inc. Google maps with streetview. `http://maps.google.com/help/maps/streetview/`, January 2010.

[8] S. Guven and S. Feiner. Visualizing and navigating complex situated hypermedia in augmented and virtual reality. In *Proceedings of the 2006 Fifth IEEE and ACM International Symposium on Mixed and Augmented Reality (ISMAR'06)-Volume 00*, pages 155–158. IEEE Computer Society Washington, DC, USA, 2006.

[9] M. Kahari and D. J. Murphy. Sensor-fusion based augmented reality with off the shelf mobile phone. In *ISMAR '08: Proceedings of the Seventh IEEE and ACM International Symposium on Mixed and Augmented Reality*, September 2008.

[10] C. Poullis and S. You. Automatic creation of massive virtual cities. In *Proceedings of the IEEE Virtual Reality Conference 2009*, pages 199–202, Lafayette, Louisiana, USA, March 2009.

[11] M. Tönnis, C. Sandor, G. Klinker, C. Lange, and H. Bubb. Experimental evaluation of an augmented reality visualization for directing a car driver's attention. In *Proc. of IEEE and ACM International Symposium on Mixed and Augmented Reality*, pages 56–59, Vienna, Austria, October 2005.

[12] S. Vallance and P. Calder. Context in 3D planar navigation. *Australian Computer Science Communications*, 23(5):93–99, 2001.