

# Continuous Navigation of Nested Abstraction Levels

Matthew van der Zwan  
 University of Groningen  
 m.a.t.van.der.zwan@rug.nl

Alexandru Telea  
 University of Groningen  
 a.c.telea@rug.nl

Tobias Isenberg  
 University of Groningen  
 DIGITEO in collaboration with VENISE–  
 LIMSI–CNRS and AVIZ–INRIA, Saclay, France  
 isenberg@cs.rug.nl

## Abstract

We investigate the dedicated control of multiple levels of semantic and sampling-based abstraction in 3D datasets, i. e., different types of data abstractions as opposed to sampling-based abstraction which shows more or less data. This dedicated navigation in the abstraction space facilitates the mental integration of different existing visualization techniques in many application areas including our example domain of fluid simulation. We realize the continuous abstraction control by interpolating between the levels while being able to simultaneously show multiple abstractions. We employ a halo-like shading technique based on distance fields to blend between several levels while continuously navigating between focus and context abstractions. We further add a semantic lens to find focus abstractions close to a user-defined context abstraction. Our entire implementation uses 2D image-based techniques to enable real-time performance, which seamlessly integrates within a 3D visualization tool.

## 1 Introduction

Abstraction is a—if not even *the*—fundamental principle employed in virtually all areas of visualization because it allows us to uncover and understand principles about the subject matter that we visualize, rather than just seeing the raw data. As Rautek et al. [21] note, abstraction can be introduced in a visualization either implicitly by selecting a certain style of depiction (“low-level visual abstractions”) or explicitly by employing means such as focus+context or distortion (“high-level visual abstractions”). The latter group of high-level abstractions are of particular interest because they are created to emphasize specific chosen aspects of interest to the viewers. Often, however, there exist many different means to achieve explicit or high-level abstractions, all of which are valid and show different important aspects of the same dataset. Therefore it is essential that we can link these different types of abstraction with each other [10] to allow viewers to understand the relationship between them.

---

This paper has been submitted for presentation at ICT.OPEN 2012. This is the national Dutch ICT conference and serves, for this paper, the role of training the presentation skills rather than to push scientific limits. ICT.OPEN does not claim copyright. For this reason, ICT.OPEN 2012 encourages authors to submit mostly papers that have been sent to or have recently been presented at international conferences.

In this paper we focus on addressing this issue of combining multiple different layers of abstraction of the same dataset specifically for 3D datasets whose abstract representations are *spatially* and *semantically* nested. By spatially nested we mean that the abstractions are defined in the same spatial embedding but each uses a different amount of screen space such that more abstract representations are, generally, ‘inside’ less abstract ones. The nesting property allows us to use two-dimensional techniques to generate real-time halos which appear volumetric and visually separate the different visualizations. This allows us to create transitions between different abstractions which do not allow seamless geometrical transitions as demonstrated using different representations of fluid flow. We also incorporate lens-based navigation into the defined abstraction space allowing investigation of a different level of abstraction. Taken together, these techniques facilitate an intuitive continuous navigation of a set of nested abstractions of a given 3D visualization.

## 2 Related Work

Abstraction is a core principle in visualization and takes many forms, depending on the visualized data. Dedicated, controlled abstraction has been investigated not only in non-photorealistic rendering (e. g., [9, 19, 28]) but also in visualization. In the field of information visualization many forms of intentional abstraction are used such as edge bundling (e. g., [15]) or focus+context (e. g., [7]). In scientific and specifically illustrative visualization many “high-level visual abstractions” [21] are used.

Relevant for our own work are those high-level visual abstractions that not only show more or less relevant parts of a dataset in more or less detail but which can relate different visual representations (i. e., different abstraction levels) to each other. Duke [10] describes this problem nicely and suggests linking different categories or representations to each other to uncover and understand the structure in a dataset, naming molecular visualization as one example. Van der Zwan et al. [26] demonstrated such seamless transition between molecular abstraction levels in an interactive [26] and a spatially explicit [18] manner. However, van der Zwan et al.’s [26] realization of abstraction level transitions requires that meaningful intermediate stages exist which is not the case for many forms of abstraction—a problem that we address in our own work.

We employ halos as one visual technique to visually separate layered elements and thus to enhance spatial perception. While halos enhance the spatial perception of the depicted objects in volume [4, 17] and other visualization domains [13, 23], we employ them to support visual layering and thus related to their function of showing occlusion relationships in line-based techniques [1, 11]. In addition, we use interactive lenses to locally explore our layered abstractions. Lenses are not only frequently used to support focus+context techniques [7] but also to interactively reveal otherwise hidden information, an approach pioneered by Bier et al. [2, 3]. Their Magic Lenses locally affect a 2D screen region using a user-selected operator. While lenses can be used in a 3D context to distort the projection [29], they can also be used to specify non-view changes for a 3D scene in a separate 2D layer [16, 20]. Our lenses have a similar function as we use them in a 2D layer over the 3D model to locally reveal relationships between abstraction layers, thus also relating to a number of smart visibility methods in visualization [27].

### 3 Visualization Model

We start with a dataset  $d \in D$  and consider several visualizations of  $d$ , modeled as images  $V_{1 \leq i \leq N} : D \rightarrow \mathbf{R}^2$  and each producing a 2D image  $A_i = V_i(d)$ . We call these images *abstractions* of  $d$  if they represent the information in  $d$  on different levels of detail. We distinguish two abstraction types: *Semantic* abstractions  $A_i$  simplify the information in  $d$  by showing varying amounts of the information present in  $d$  using different visual representations. For example, a fluid flow volume  $d \subset \mathbf{R}^3$  can be rendered as an entire flow volume using LIC [5, 22], as stream LIC structures for a set of given streamlines [14], and as flow topology [25]; these are increasingly simplified semantic representations. *Sampling* abstractions reduce the amount of points produced by a given semantic abstraction  $A_i$  using data sampling. Rendering different numbers of streamlines, for example, are samplings of the streamline abstraction. We denote all  $S_j$  samplings of a semantic abstraction  $A_i$  by  $A_i^j$ ,  $1 \leq j \leq S_i$  with  $A_i^1 = A_i$  the most detailed and  $A_i^{S_i}$  the coarsest sampling.

To be useful in an exploration scenario, abstractions must be described in terms of the amount of *simplification* they produce on some input dataset. In our model we assume that, for a given application domain (e. g., flow visualization), the abstraction set  $A = \{A_i\}$  can be ordered in decreasing amount of provided simplification from the densest abstraction  $A_1$  to the sparsest one  $A_N$ . We also require that simpler abstractions are visually *nested* within less simple ones, i. e.  $A_j \subset A_i, \forall i < j$ . This is often the case in scientific visualization where abstraction reduces the size and/or spatial dimensionality of the depicted visual elements while keeping them aligned in the space  $D$ . Our flow visualization scenario is such a case of nested abstractions: the topology is a part of the streamlines, these are nested within stream LIC representation, which in turn is a part of a LIC volume.

### 4 Navigating the Abstraction Space

Given an abstraction set  $A$  as just described, one typically wants to *navigate*  $A$  to get different types of insight which are best visible at different abstraction levels. One navigation option is to start with  $A_N$  (most abstract) and browse through  $A_i$  until  $A_1$  (most detailed), optionally using spatial sampling to restrict the dense-data areas to zones of interest using, e. g., focus+context techniques. One can also start with the most detailed level  $A_1$  and simplify the visualization to the coarsest level  $A_N$  is reached. During both navigation types, we call the level of the highest abstraction  $A_f$  being visualized the *focus* of the visualization: Given a user-selected  $f$ , we aim to produce a visualization combining all *context* abstractions  $A_i, i < f$  and  $A_f$  in a single visualization such that all abstractions and their spatial nestings are shown. This will permit smooth navigation in the *combined* space of semantic and sampling abstractions (as introduced in Section 3).

Such navigations are typically realized by toggling the rendering of the elements  $A_i$  on and off. However, this creates sharp visual discontinuities in the transition, especially if the abstractions differ visually. Continuity can be added by smoothly interpolating the transparency or shape of consecutive  $A_i$  using fading or morphing while navigating through  $A$ . However, blending blurs the spatial nesting insight and can result in too high opacities when too many abstractions are blended. Morphing is not trivial for any pair of (nested) shapes, works only for shape pairs, and requires 3D shape representations rather than their 2D visualization results  $A_i$ .

We propose to create a continuous navigation function  $Nav : A \times [0, 1] \rightarrow \mathbf{R}^2$  to help navigation in the abstraction space. Given our set  $A$  of ordered, nested abstractions and a focus abstraction level  $f \in [0, 1]$ , we combine all abstractions  $A$  to build a visualization  $V$ . As the user changes the focus level  $f$ ,  $V$  continuously changes to show only  $A_1$  (at  $f = 0$ ), next show the focus abstraction  $A_f$  nested within lower abstractions as context, and finally  $A_N$  (at  $f = 1$ ). The design of  $Nav$  should be such that it can be computed on any set of nested 2D or 3D abstractions, is continuous in  $f$ , clearly emphasizes the focus-context relation of nested abstractions, and is computed using only 2D image information instead of 3D shape information to achieve maximal performance.

We use an additive blending of the abstractions  $A_i$  in nesting levels (decreasing  $i$ ) and compute the navigation function as  $Nav(A, f) = \sum_{i=1}^n \alpha_i(f) A_i$  (see Fig. 1). The design of the blending factors  $\alpha_i : [0, 1] \rightarrow [0, 1]$  is essential, we use  $\alpha_i(f) = \phi_i(f) \cdot \psi_i(f) \cdot h_i(f)$ . Here,  $\phi_i : [0, 1] \rightarrow [0, 1]$  is the function used to fade in an abstraction  $A_i$ ,  $\psi_j : [0, 1] \rightarrow [0, 1]$  is the function responsible for fading out a (context) abstraction  $A_j, j < i$ , and  $h_i : [0, 1] \rightarrow [0, 1]$  is the *halo function* used to create a halo around the selected abstractions. As fade-in function we use  $\phi_i = \max\left(0, \min\left(1, \frac{f - f_{in}^{(i)}}{f_{full}^{(i)} - f_{in}^{(i)}}\right)\right)$  where  $f_{in}^{(i)}$  is the focus value from which we start fading in abstraction  $A_i$  and  $f_{full}^{(i)}$  is the value at which  $A_i$  is completely visible. In practice, we want to start fading in the next abstraction  $A_{i+1}$




Figure 1: Continuous navigation in a flow visualization abstraction space with four abstractions  $A_1$ – $A_4$ .




Figure 2: Construction of the focus guided lens. Typical values are  $\rho_c = 5$  pixels and  $\rho_{A_i(c)} = 90$  pixels.

when the current abstraction  $A_i$  is fully visible. Hence, we choose  $f_{in}^{(i+1)} = f_{full}^{(i)}$ . Similarly, we define the fade-out function  $\psi_i = \min\left(1, \max\left(0, \frac{f_{gone}^{(i)} - f}{f_{gone}^{(i)} - f_{out}^{(i)}}\right)\right)$ . Here,  $f_{out}^{(i)}$  is the value for which we start fading out the context abstraction  $A_i$  and  $f_{gone}^{(i)}$  is the value for which abstraction  $A_i$  is no longer visible. Similar to fading in, we start fading out abstraction  $A_{i+1}$  when the  $A_i$  is no longer visible and, thus, set  $f_{gone} = f_{out}^{(i+1)}$ . Also, we constrain  $f_{out}^i > f_{full}^{i+1}$  such that  $A_i$  does not start fading out before  $A_{i+1}$  is fully in focus. Finally, we set  $f_{full}^1 = 0$ ,  $f_{out}^N = 1$  so that we start with a fully focused  $A_1$  and end with a fully focused  $A_N$ .

While combining  $\phi_i$  and  $\psi_i$  allows us to fade abstractions in and out of view continuously, the resulting image will be unable to clearly show the *nesting* structure of the abstraction space: Depending on the specific abstraction image shapes and colors, it may be hard to see which result pixels belong to a (thin) abstraction being nested within a (larger) context abstraction, especially if both have similar colors. We therefore use the *halo function*  $h_i$  to generate halos around abstractions:  $h_i(f) = \min\left(\left(DT_{A_{i+1}}/\delta\right)^{k_i(f)}, 1\right)$ . In this function,  $DT_\Omega : \mathbf{R}^2 \rightarrow \mathbf{R}^+$  is the distance transform of a 2D binary shape  $\Omega$ , which gives, for any points  $\mathbf{x} \in \mathbf{R}^2$ , its distance to  $\Omega$  [8].  $DT$  is zero inside  $\Omega$  and smoothly increases outside the shape. In our case, we construct such shapes by sim-

ply thresholding the rendered abstractions  $A_i$  into foreground (rendered) and background (not rendered) pixels. Having  $DT_{A_{i+1}}$ , we compute a halo around  $A_{i+1}$  by modulating the distance transform with a power function  $k_i(f)$ . The halo's width is limited to a maximal value of  $\delta > 0$  pixels. The effect of the power function is to create a smoother transition from the context  $A_i$  of  $A_{i+1}$  than if linear distance functions were used. Finally, we set  $k_i = \phi_i$  to increase the halo around the fading-in abstraction  $A_{i+1}$ , thus making it more prominent in its context  $A_i$  where it is nested. Perspective-like halos can easily be obtained by modulating the value of  $\delta$  with the depth of  $A_{i+1}$  at each pixel.

The above process describes how *semantic* abstractions  $A_i$  are combined into a single image. However, as defined in Section 3, our input may contain sampled versions thereof. We integrate these smoothly in the above process by replacing, in the navigation function  $Nav(A, f)$ , each semantic abstraction  $A_i$  with its sampled version  $A_i^j$ , the sampling parameter  $j$  being controlled by the distance from the user-set focus  $f$  to the full-visibility  $f_{full}^i$  as  $j = \frac{f - f_{full}^i}{f_{out}^i - f_{full}^i} S_i$ . In other words, as the user increases  $f$ , the full-visibility abstraction is progressively simplified from  $A_i^1$  to  $A_i^{S_i}$  (coarsest variant). During this process, all remaining visualization elements stay the same (halo sizes, overall abstraction transparency). When  $f$  reaches  $f_{out}^i$ , the coarse abstraction  $A_i^{S_i}$  is further faded out. For abstractions which have no level-of-detail representations, the process simply uses the unique representation  $A_i$ . This directly accommodates any number of semantic abstractions with any number of sampling representations thereof, effectively intertwining the navigation in the semantic and sampled spaces of abstractions.

## 5 Interactive Local Exploration

While this navigation facilitates an effective *global* abstraction space exploration, in many cases we are interested in getting local detail information. We thus also provide context-sensitive *local* lenses, whose goal it is to allow parts of an abstraction  $A_{l>f}$  located inside the lens to become visible even when otherwise hidden due to the global abstraction level  $f$ . While a naïve implementation would change the blending factors  $\alpha_i(f)$  close to the lens center, this would interfere with our transparent distance-based halos. In multi-layer visualizations such as ours one also wants to see ‘deeper’ within the abstraction stack inside the lens *and* locate the parts of deeper-nested abstractions closest to the lens.

For this, we use a *focus-guided lens*. Given a global abstraction level  $f$ , we combine revealing deeper-nested information at the lens center with revealing higher-abstraction structures  $A_{l>f}$  close to it. We first locate the closest point  $A_l(c)$  of abstraction  $A_l$  to the lens center  $\mathbf{c}$ . The point  $\mathbf{c}$  can be directly computed as  $A_l(c) = FT_{A_l}(\mathbf{c})$ . Here,  $FT_{A_l} : \mathbf{R}^2 \rightarrow \mathbf{R}^2$  is the *feature transform* of the shape  $A_l$  [8]. The feature transform of a shape  $\Omega \subset \mathbf{R}^2$  is defined as  $FT_\Omega(\mathbf{x}) = \{\mathbf{y} \in \Omega | DT_\Omega(\mathbf{x}) = \|\mathbf{x} - \mathbf{y}\|\}$ , i.e. the closest point  $\mathbf{y} \in \Omega$  to a given target point  $\mathbf{x}$ , restricting ourselves to a one-point feature




Figure 3: Navigating the flow visualization abstraction space: (a)–(c) Introducing seed LIC as focus in the LIC volume, (d) has the flow topology as focus abstraction and the previous abstractions as context from which the LIC volume is removed in (e).




Figure 4: Fluid flow abstracted (a)–(b) locally with a guided lens and (c)–(d) with sampling abstractions (streamline filtering).

transform [12]. With the lens center  $\mathbf{c}$  and closest abstraction point  $A_l(\mathbf{c})$ , we construct our focus-guided lens by multiplying the halo functions  $h_{l>f}$  with the distance transform of a beam-like shape created by two circles connected by a trapezium (light blue in Fig. 2). As the lens is moved, it behaves similarly to a light beam that shows the shortest spatial path from the lens center to the desired  $A_l$ . This is useful as one does not need to fully remove (make transparent) *all* abstractions  $A_{k<l}$  in order to discover  $A_l$ . Hence, one can stay at a desired semantic focus level  $f$  and use the lens to search for another desired  $A_{l>f}$  in the vicinity of any point.

### 6 Implementation and Results

For the realization we only require a set of  $N - 1$  2D images depicting the different context abstraction levels and one image depicting the abstraction in focus at the selected abstraction level, as our method works entirely in image space. These images are either generated on-the-fly or are precomputed if they do not change during the exploration. From these images we compute the soft halos for our nesting (within 10 ms on a modern graphics card) by employing a recent CUDA-based implementation [24] of exact Euclidean distance transforms and feature transforms [6]. Finally, blending is implemented via OpenGL alpha blending. The entire process, including rendering LIC, seed LIC, streamlines, and precomputed topology, works at 5 frames per second on a MacBook with 2 GB RAM and an NVIDIA GeForce 320M graphics card with 256 MB RAM. A video illustrating our method is available at <http://tobias.isenberg.cc/VideosAndDemos/Zwan2012CNN>.

Fig. 3 shows our global abstraction applied to a direct incompressible Navier-Stokes flow simulation around a cylinder. Fig. 3(a)–(c) show three transition steps of stream LIC blended with the LIC volume separated by the halo around

the new focus, *i.e.* a visualization with the streamline LIC as focus and the remainder (LIC) as context. A further abstraction (Fig. 3(d)) adds topology as the new focus. Finally, we remove the lowest abstraction (Fig. 3(e)). Our second example (Fig. 4) shows the local exploration of a different Navier-Stokes simulation, with Fig. 4(a) and (b) showing two lens locations. Fig. 4(c) and (d), finally demonstrate the inclusion of several sampling abstraction levels into the treatment.

### 7 Conclusion

In this paper, we have presented an abstraction technique for continuous navigation of nested abstraction levels. This navigation can be achieved on both global and local levels, using lenses for the latter. For both types of navigation, we use two-dimensional distance transforms to create smooth halos in the image domain which look like three-dimensional volumetric halos. As an example application, we used our technique to navigate different semantic and sampling abstractions of fluid flow visualization.

### References

- [1] Arthur Appel, F. James Rohlf, and Arthus J. Stein. The Haloed Line Effect for Hidden Line Elimination. *ACM SIGGRAPH Computer Graphics*, 13(3):151–157, August 1979.
- [2] Eric Bier, Maureen Stone, and Ken Pier. Enhanced Illustration Using Magic Lens Filters. *IEEE Computer Graphics and Applications*, 17(6):62–70, November/December 1997.
- [3] Eric A. Bier, Maureen C. Stone, Ken Pier, William Buxton, and Tony D. DeRose. Toolglass and Magic Lenses: The See-Through Interface. In *Proc. SIGGRAPH*, pages 73–80, New York, 1993. ACM.

- [4] Stefan Bruckner and Eduard Gröller. Enhancing Depth-Perception with Flexible Volumetric Halos. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1344–1351, November/December 2007.
- [5] Brian Cabral and Leith Casey Leedom. Imaging Vector Fields using Line Integral Convolution. In *Proc. SIGGRAPH*, pages 263–270, New York, 1993. ACM.
- [6] Thanh-Tung Cao, Ke Tang, Anis Mohamed, and Tiow-Seng Tan. Parallel Banding Algorithm to Compute Exact Distance Transform with the GPU. In *Proc. 13D*, pages 134–141, New York, 2010. ACM.
- [7] M. S. T. Carpendale and Catherine Montagnese. A Framework for Unifying Presentation Space. In *Proc. UIST*, pages 61–70, New York, 2001. ACM.
- [8] L. Costa and R. Cesar. *Shape Analysis and Classification: Theory and Practice*. CRC Press, 2000.
- [9] Doug DeCarlo and Anthony Santella. Stylization and Abstraction of Photographs. *ACM Transactions on Graphics*, 21(3):769–776, July 2002.
- [10] David J. Duke. Linking Representation with Meaning. In *Posters of IEEE Visualization*, Los Alamitos, 2004. IEEE Computer Society.
- [11] Gershon Elber. Line Illustrations  $\in$  Computer Graphics. *The Visual Computer*, 11(6):290–296, June 1995.
- [12] O. Ersoy, C. Hurter, F. Paulovich, G. Cantareira, and A. Telea. Skeleton-Based Edge Bundling for Graph Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2364–2373, December 2011.
- [13] Maarten H. Everts, Henk Bekker, Jos B. T. M. Roerdink, and Tobias Isenberg. Depth-Dependent Halos: Illustrative Rendering of Dense Line Data. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1299–1306, November/December 2009.
- [14] A. Helgeland and O. Andreassen. Visualization of Vector Fields using Seed LIC and Volume Rendering. *IEEE Transactions on Visualization and Computer Graphics*, 10(6):673–682, November/December 2004.
- [15] Danny Holten. Hierarchical Edge Bundles: Visualization of Adjacency Relations in Hierarchical Data. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):741–748, September/October 2006.
- [16] Christophe Hurter, Alexandru Telea, and Ozan Ersoy. MoleView: An Attribute and Structure-Based Semantic Lens for Large Element-Based Plots. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2600–2609, December 2011.
- [17] Victoria Interrante and Chester Grosch. Visualizing 3D Flow. *IEEE Computer Graphics & Applications*, 18(4):49–53, July 1998.
- [18] Wouter Lueks, Ivan Viola, Matthew van der Zwan, Henk Bekker, and Tobias Isenberg. Spatially Continuous Change of Abstraction in Molecular Visualization. In *Abstracts of IEEE BioVis*, Los Alamitos, 2011. IEEE Computer Society.
- [19] Xiaofeng Mi, Doug DeCarlo, and Matthew Stone. Abstraction of 2D Shapes in Terms of Parts. In *Proc. NPAR*, pages 15–24, New York, 2009. ACM.
- [20] Petra Neumann, Tobias Isenberg, and Sheelagh Carpendale. NPR Lenses: Interactive Tools for Non-Photorealistic Line Drawings. In *Proc. Smart Graphics*, pages 10–22, Berlin, Heidelberg, 2007. Springer-Verlag.
- [21] Peter Rautek, Stefan Bruckner, Eduard Gröller, and Ivan Viola. Illustrative Visualization: New Technology or Useless Tautology? *ACM SIGGRAPH Computer Graphics*, 42(3):4:1–4:8, August 2008.
- [22] Detlev Stalling and Hans-Christian Hege. Fast and Resolution Independent Line Integral Convolution. In *Proc. SIGGRAPH*, pages 249–256, New York, 1995. ACM.
- [23] Marco Tarini, Paolo Cignoni, and Claudio Montani. Ambient Occlusion and Edge Cueing to Enhance Real Time Molecular Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):1237–884, September/October 2006.
- [24] A. Telea. CUDA Skeletonization and Distance Transform Toolkit, 2011. <http://www.cs.rug.nl/~alex/CUDASKEL>.
- [25] Holger Theisel, Tino Weinkauff, Hans-Christian Hege, and Hans-Peter Seidel. Saddle Connectors – An Approach to Visualizing the Topological Skeleton of Complex 3D Vector Fields. In *Proc. IEEE Visualization*, pages 225–232, Los Alamitos, 2003. IEEE Computer Society.
- [26] Matthew van der Zwan, Wouter Lueks, Henk Bekker, and Tobias Isenberg. Illustrative Molecular Visualization with Continuous Abstraction. *Computer Graphics Forum*, 30(3):683–690, May 2011.
- [27] Ivan Viola and Eduard Gröller. Smart Visibility in Visualization. In *Proc. CAE*, pages 209–216, Goslar, Germany, 2005. Eurographics Association.
- [28] Georges Winkenbach and David H. Salesin. Computer-Generated Pen-and-Ink Illustration. In *Proc. SIGGRAPH*, pages 91–100, New York, 1994. ACM.
- [29] Y. Yang, J. Chen, and M. Beheshti. Nonlinear Perspective Projections and Magic Lenses: 3D View Deformation. *IEEE Computer Graphics & Applications*, 25(1):567–582, January/February 2005.