

Integrating Expanding Annotations with a 3D Explosion Probe

Henry Sonnet
Dept. of Simulation and
Graphics
University of Magdeburg
Universitaetsplatz 2
D-39106 Magdeburg,
Germany
sonnet@isg.cs.uni-
magdeburg.de

Sheelagh Carpendale
Dept. of Computer Science
University of Calgary
2500 University Drive N.W.
Calgary, Alberta,
Canada T2N 1N4
sheelagh@cpsc.ucalgary.ca

Thomas Strothotte
Dept. of Simulation and
Graphics
University of Magdeburg
Universitaetsplatz 2
D-39106 Magdeburg,
Germany
tstr@isg.cs.uni-
magdeburg.de




Figure 1: (a) Anatomical diagram of a foot and (b) explosion diagram with an interior bone labeled.

ABSTRACT

Understanding complex 3D virtual models can be difficult, especially when the model has interior components not initially visible and ancillary text. We describe new techniques for the interactive exploration of 3D models. Specifically, in addition to traditional viewing operations, we present new text extrusion techniques combined with techniques that create an interactive explosion diagram. In our approach, scrollable text annotations that are associated with the various parts of the model can be revealed dynamically, either in part or in full, by moving the mouse cursor within annotation trigger areas. Strong visual connections between model parts and the associated text are included in order to aid comprehension. Furthermore, the model parts can be separated, creating interactive explosion diagrams. Using a 3D probe, occluding objects can be interactively moved

apart and then returned to their initial locations. Displayed annotations are kept readable despite model manipulations. Hence, our techniques provide textual context within the spatial context of the 3D model.

Categories and Subject Descriptors

I.3.6 [Methodology and Techniques]: Interaction techniques; H.4 [Information Systems Applications]: Miscellaneous; H.5.2 [User Interfaces]: Graphical User Interfaces (GUI); I.3.8 [Computer Graphics]: Applications.

Keywords

3D model exploration, expanding annotations, explosion diagram, interaction design.

1. INTRODUCTION

Powerful graphics computers are now widely available and can provide environments for exploring virtual 3D models. These models can represent a great variety of information from many domains. Exploring a 3D model of an engine, for example, can help students, studying auto mechanics, learn about the engine components and how they are assembled. In the medical domain, 3D models of human organs are already used for diagnostic purposes and for the planning of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
AVI '04, May 25-28, 2004, Gallipoli (LE), Italy
© 2004 ACM 1-58113-867-9/04/0500..\$5.00

surgical procedures. Text annotations can be attached to specific model components to facilitate the exchange of diagnostic results between surgeons.

When compared to printed media, interactive environments for exploring virtual 3D models offer several advantages. In particular, they permit more freedom of exploration. The user can choose arbitrary views of the model and decide when and which pieces of additional information are shown. Although the focus of this discussion concerns additional textual information, similar claims can be made for multimedia information such as movies or sound clips. However, most of the methods available for model exploration are limited to simple camera rotations and annotations with pre-defined visibility and placement.

By examining a 3D model, the user can gather information about color, surface quality, shape, and relationships in terms of scale, size, topology and inter-connections of the different parts. Typically, this is the information that has been used to create the 3D model and as a result, this information is well integrated into the model. Often these model features alone do not fully explain the 3D model; therefore, to provide more complete explanations, additional text information is associated with the model. It is still common for the related textual information to be made available separately. This separation can vary from the extreme of off-line text, to separate text windows, to integrated labels that refer to separate text. When the additional text information is presented with the model, it is most often shown in a predetermined layout and can only be displayed or hidden as it was designed. Also, when the text is shown with the model, in order to fit it into limited space and avoid occlusion of the 3D model, the textual information is usually quite abbreviated.

Both the visual model information and the textual information provide their own type of context. Our research challenge has been to dynamically integrate these two types of information, providing exploration tools that maintain both textual and spatial context (see Figure 1). In order to explore visual model information, while maintaining the appropriate shapes and relative spatial relations with nearby components, we adapt a radial explosion technique (see CARPENDALE et al. [2]). Furthermore, with our exploration tools, a person is in control of which model parts will display their associated textual information, how much of this text will be displayed, and whether this textual context will persist through the subsequent exploration of the model.

This paper is organized as follows: After reviewing related work in Section 2, we discuss the general idea behind interactively exploring 3D models in Section 3. The spatial exploration of the model using a 3D probe is introduced in Section 4. In Section 5 we show how distortion lenses can be used to control the size, shape, and content of annotations. Discussion and future directions are the contents of Section 6 and Section 7 summarizes this paper.

2. RELATED WORK

2.1 Concept of Context

Our research applies the basic ideas behind *focus plus context* research (e.g., [8, 18]) to the maintenance of both spatial and textual context. *Focus plus context* research has generated many methods by which selected aspects, or foci,

of a presentation can be enlarged or magnified without removing them from their context (e.g., [15, 17]). Magnified foci are usually achieved by compressing the context just enough to create the necessary space for the enlargement. In contrast, we do not use any local magnification; however, we do use the concept of radially diminishing functions that is common to many *focus plus context* techniques. We apply these drop-off functions to foci, which can be used to trigger our expanding labels or to compute focus-dependent explosion diagrams.

2.2 Integration of Text and 3D Presentation

Our technique for dynamically accessing text labels is related to a small but growing body of research that is interested in integrating text with visual representations. Since the exploration of 3D models is intimately associated with showing textual information, some elementary rules concerning annotations should be followed. For example, labels should always be readable, not hide other information, and be non-ambiguously related to their counterparts [7]. In this regard, BELL et al. [1] use the upright rectangular extent of object projections on the view plane to determine where text labels should be placed. In an application called *Illustrative Shadows* [16], the shadow of 3D objects is utilized for both emphasizing scene objects and as reference for annotations. CHIGONA et al. [5] also use the shadow for illustration. In their approach, text is placed within the shadow of the corresponding objects. However, the shape of objects may limit the amount of information to be shown as well as its readability.

In most applications that support the exploration of 3D models, the text and the model are still quite separate. For example, JUNG et al. [9] have developed a system for collaborative work that allows people to leave notes symbolized by geometric objects integrated with the 3D model. The textual content of those geometric objects (the notes) can be retrieved in a separate window. Separating text and model is suited for dividing the two tasks of reading the text and exploring the visual information of the model. However, the distance between text and model is fairly large and there is no direct correlation between them. In *Voxel-Man* [13], labels are connected to the appropriate objects with lines. Other information about the objects can be accessed through pop-up menus. But these menus are not visually connected to their objects. More extended textual information is shown in a separate window. PREIM et al. [14] place annotation boxes at fixed positions beside the model. These annotations are connected to the model parts with referring lines as well. However, these lines can become confusing, especially after changing the viewpoint.

2.3 Interactive 3D Environments for Spatially Exploring Models

As MCGUFFIN et al. [11] have observed, methods for exploring 3D models usually include: common camera transformations, various filtering techniques such as transparency and cutting planes (as in [13]), and relatively few alternative methods that usually use some type of deformation or spatial transformation. Our explosion technique belongs in the last group.

A few deformation techniques have been developed that show the interior of volumetric data. KURZION and YAGEL [10] use a *ray deflector* to simulate a deformation of model

parts. CARPENDALE et al. [3] use a book metaphor to create an effect similar to opening between pages. A variety of 3D widgets for exploring the interior of volume data while retaining surrounding context has been introduced by MCGUFFIN et al. [11]. They provide tools for spreading volumes, expanding volumes from inside, or hinging open parts of a volume. *3D Magic Lenses* [19] are tools that can display alternative information within a chosen region of a 3D model. RAAB and RÜGER [15] have developed a 3D fish-eye zoom technique that enables the enlargement of selected objects in a 3D model while surrounding objects are scaled down. This technique is utilized by PREIM et al. [14] to emphasize objects that are currently in focus and to control the layout of annotation boxes.

Our research relates more closely to the concept of explosion diagrams [12] and the concept of radial expansion provided by *Visual Access Distortion* as outlined in [2]. We apply *Visual Access Distortion* to 3D models that contain component parts of many different shapes and sizes.

3. EXPLORING 3D MODELS

A 3D model has a visual spatial representation and may be accompanied by extensive textual explanations. The visual representation contains information, much of which is inherent in its spatial organization and construction. The amount of associated textual explanations can vary from a few labels to many pages. Gaining understanding and insight from a 3D model requires exploring both the model and text, and integrating the verbal information with the visual representation. Our goal is to create interactive tools that will provide support for this cognitive task. To do this we:

- create explosion diagrams that operate interactively with 3D models; and,
- allow the user to request accompanying text for model components that require explanation.

The text appears on demand. It is positioned so as to be legible and to clearly indicate the model part with which it is associated. Also on demand, the text area can expand sufficiently to present more readable text. We first describe the explosion diagrams before the expanding labels are introduced.

4. EXPLOSION DIAGRAMS




Figure 2: (a) 3D models of a heart and (b) an engine.

4.1 Challenges of 3D Model Exploration

A 3D model is usually arranged in a pre-defined constellation. The constellation is based on the relationships between and the functionality of the model's components. This spatial organization is often an important part of the information available through exploring the model. However, this same organization has objects of many different shapes and sizes and often contains interior and occluded objects. Figure 2, for instance, shows two relatively simple models, yet both these models contain interior and occluded or partially occluded parts. In order to explore interior objects of a 3D model, occluding objects have to be removed in some way. This can be done by cutting sections away, by rendering parts transparent, by using deformations, or by moving objects aside (see Section 2.3). Cutting sections away will reveal interior regions; however, a part of the context has been removed. Rendering some parts transparent does a better job of maintaining context but provides partially obscured views. Deformations provide new exploration metaphors such as explosion diagrams and leafing through volume data. Moving objects aside has promise, but if this is done by selecting and transforming each object individually it has several disadvantages:

- The interaction effort increases as the number of occluding objects to be moved increases.
- If user interaction is depended on to replace objects, this task can also be tedious and difficult to be precise.
- Choosing where to place displaced objects can be difficult.

This last point is especially important because careful placement of the displaced objects can retain or even enhance model construction information (see NIEDERAUER et al. [12]).

4.2 Interactive Approach

To provide visual access to interior objects of the model—without distorting the objects' shapes or losing spatial context—we chose to extend *Visual Access Distortion* (see CARPENDALE et al. [2]). *Visual Access Distortion* operates according to a line of sight between an object of interest and the viewpoint. It does not necessarily apply magnification or compression. It merely notes which objects obstruct the line of sight and moves them out of the way to clear the line of sight. We make two extensions to this technique. *Visual Access Distortion* has been applied to node and edge graphs in which all the objects are small cubes. In contrast, the objects of 3D models can be any type of shape or size. To handle these different objects we make use of bounding boxes, skeletons, and centroids. Also, in *Visual Access Distortion* the line of sight is terminated at a data item, treating this data item as a focus. We remove this restriction and allow the terminus of the line of sight to be any spatial location. Combining this with the ability to clear objects from the line of sight creates a *3D probe* that can be used to interactively explore a 3D model.

4.3 3D Probe

The purpose of a 3D probe is to provide the user with an interactive widget that temporarily modifies the 3D model, making occluded objects visible. Instead of selecting every single scene object and repositioning it, the user can move

the 3D probe. As the probe is moved into the model, the line of sight between the probe and the user is kept clear. Thus, the probe has the effect of pushing the model components apart, revealing chosen inner reaches. Objects that are close to the probe's line of sight are moved apart radially. The distance they are moved depends upon their initial distance to the line of sight, their size, and the probe's chosen drop-off function. As soon as objects leave the scope of the probe, they are moved back to their initial location.

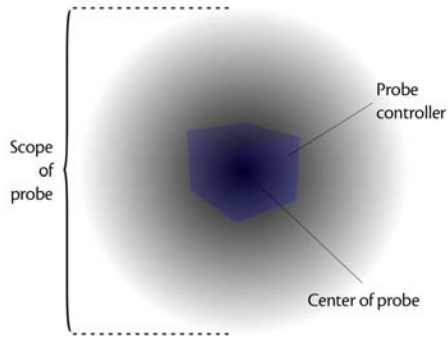


Figure 3: Sketch of a probe's functional components. The centre of the probe is the centre of the cube. The cube is what the user will see and use to position the probe. The scope of the probe is indicated by the transparent circle.

Figure 3 illustrates both the visible and invisible parts of the probe. The cube is what the user will see and use to position the probe. To move the cube, the user indicates one of the cube's faces and then, the cube will move with mouse repositioning, in the plane associated with the chosen face. The centre of the cube is the centre of the probe and the terminus of the line of sight. The probe has a degree of influence that is at maximum along the line of sight to the probe centre and diminishes radially, defining the probe's scope. The circular shape in Figure 3 indicates the probe's scope. Objects outside the probe's scope are not moved.

We adapted the scope of the probe to the size of the objects in order to avoid introducing strong effects on small objects or, nearly no effect on huge objects. Within a pre-processing step, for each object a value S_O that relates to the size of the object's bounding box is calculated. During the exploration of the model, the scope of the probe is adapted according to the current object's S_O . This is done for every scene object before the respective displacement is calculated. This way, small objects that are enclosed in larger objects are moved less than the larger objects and thus, become visible. However, there is still an open problem with objects that include other objects of similar size because all objects are moved to the same position such that inner objects remain hidden.

4.4 Moving Scene Objects Apart

In order to move scene objects apart, one point of each object has to be chosen that will be affected by the probe. Once the new coordinate location for this representative point has been calculated according to the distortion caused by the probe, the difference to the initial position is applied to all other vertices of the same object. As the representative point, we have chosen the centroid of each object (see

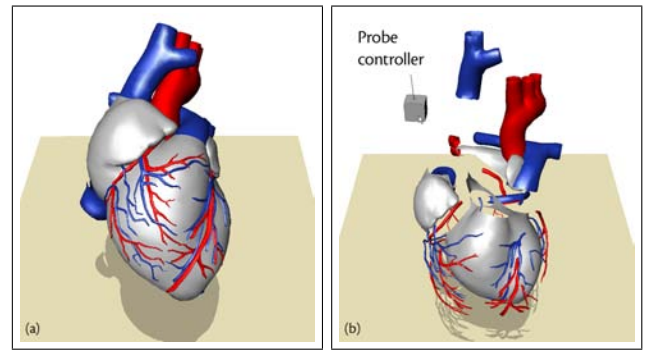


Figure 4: Effect of a 3D probe. Objects are displaced when they are covered by the probe's scope.

Section 4.5). This position is calculated once during pre-processing.

Figure 4(b) illustrates the exploration of a human heart using this technique. The probe is moved towards the upper region of the heart. All objects that are located within the probe's scope are moved apart. Hence, the user is able to see objects that were previously occluded.

4.5 Calculation of Centroids

In order to spatially explore a 3D scene by utilizing a probe, a position (the centroid) of each object has to be determined that is located inside the object and close to the object's center position. A simple way of determining the center position of an object is to calculate the center of the object's bounding box. However, the resulting position might be located outside of the object.




Figure 5: (a) 3D object whose skeleton is used to determine a significant point. Usually this point is located where most of the line segments meet each other. But this location is not necessarily the center of the object. Hence, the center of the object's bounding box and the skeleton of the object are utilized to determine the centroid of the object (b).

Another approach is based on the skeleton of an object. Assuming that the model consists of triangles, the skeleton can be computed using edge collapse operations, which yields a decreasing number of edges. This is done by merging the two vertices of an edge into one new vertex. In order to compute the skeleton, edge collapse operations are applied until no more regular triangles exist (see DEUSSEN et al. [6]). After the skeleton has been computed, one significant point that is part of the skeleton is determined. In branched skeletons a significant point is usually a point at which most of the skeleton segments coincide (see Figure 5(a)). Hence, a significant point is an end point that several line segments have

in common. If the skeleton has no branchings, the center of the skeleton is considered as the significant point.

As can be seen in Figure 5(a), the significant point is not necessarily located at the center of the object. In order to improve our method, the bounding box is incorporated into the process of finding the centroid. To this end, the point of the skeleton that is closest to the center of the bounding box is considered as the centroid of the object (see Figure 5(b)).

5. EXPANDING LABELS PROVIDING ACCESS TO ANNOTATIONS

In this section, we introduce a methodology to present text information adjacent to associated scene objects. With this technique, the user is able to influence the degree to which the annotation is revealed and the shape of the text region. Visually, annotations are directly connected to their associated objects. Though these expanding annotations exhibit no distortion, we have used a distortion algorithm to create a non-linear mapping between the spatial movement the user makes with the cursor and the amount of textual information that is revealed. Thus, approaching an objects' centroid expands the associated annotation to provide more of the text explanation.

5.1 Distortion Algorithm Basics

Distortion lenses can be considered as parameterizable objects. The basic components of a distortion lens are:

- the focal point, or region, at which the distortion is maximal,
- the area, or scope of influence, in which the lens creates an effect, and
- the manner in which this effect attenuates (drop-off function).

A Gaussian drop-off function, for example, provides a smooth transition between intensive and marginal distortions. Many drop-off variations are possible by modifying the parameters of a chosen function or by choosing a different function (for further discussion see [4]). A lens has influence only within its surrounding area or scope. The diameter of the lens' scope is adjustable. In addition, the size of region in which maximum distortion occurs can be changed at will. It can be increased, for example, by changing a focal point to a circular focal region and increasing its diameter.

5.2 Correlation between the Lens and the Text Navigation

Every scene object is associated with an object-attached lens. This object-attached lens has three parts. The first is an attenuated field that is centered at the object's centroid. The second is the expanding region that displays the text of the annotation and the third is the connection between the expanding region and the associated scene object, which is drawn to create a visual link. The attenuated field, placed at the object's centroid, is created through use of a distortion algorithm. Since the centroid of an object does not change, it is determined in a preprocessing step (see Section 4.5). This position, projected onto the viewplane, is used for placing the lens. Since the projection depends on the camera position and orientation, it is not done until the user explicitly indicates interest in a specific object. When an object is

selected, the cursor position relative to the object-attached lens center yields a distortion value. This value is used to calculate the boundary vertices of the annotation display region and thus, it affects the amount of information that is shown. As the cursor approaches the centroid, the annotation region expands, reaching a maximum when the cursor is located at the centroid. Conversely, the magnification decreases as the cursor is moved away from the centroid. The rate of expansion/contraction depends on the chosen drop-off function.

It may happen that the available text is too long for the annotation area. In this case, the user can scroll through the text. Similarly to controlling the annotation area by moving the mouse cursor in correspondence to the object's centroid, scrolling through the text is also affected by the determined distortion. But in contrast to showing the largest possible size of the annotation at maximum distortion, the text is scrolled to its beginning. Accordingly, a decreasing distortion causes the text to scroll to the end.

5.3 Determining the Appearance of the Displayed Annotation

Shapes of annotations

The text is rendered within a polygonal shape. The coordinates of this shape are affected by the distortion of the lens and the 2D bounding box of the associated 3D object projected onto the view plane. Two adjacent corner vertices of this bounding box (edge E) are selected as initial coordinates of the annotation shape. Edge E is chosen as follows:

- Edge E is one of the two longest edges of the 2D rectangular bounding box.
- Edge E is located most closely to the boundary of the view plane such that the annotation is scaled up—with an increasing distortion—towards the boundary of the view plane.

Figure 6 illustrates three different shapes of annotations. Using the determined edge E , two other vertices are calculated according to the distortion such that E and the two vertices form a rectangle R . The annotation shapes in Figures 6(a) and 6(b) have R as their bounding box. In contrast, the annotation of Figure 6(c) is directly connected to the corresponding 3D object. For this annotation, the determined edge E is used to find vertices of the projected 3D object that can be used as part of the annotation shape. In order to find suitable object vertices, edge E is evenly subdivided. Thereafter, for each segment of the subdivided edge, vertices of the projected 3D object are searched to find the vertices that have the shortest distance to E . Having determined these vertices, the two remaining vertices of the annotation shape at the opposite side are calculated according to the distortion.

Since the shape of the annotation in Figure 6(c) is partly adapted to the corresponding 3D object, the connection between annotation and object is more obvious to the viewer. However, the regularly placed text within the rectangular shaped annotation enhances the readability and thus, is best suited for longer textual information. The annotation, shaped as an ellipse, is comparable to speech bubbles that are used in comics.

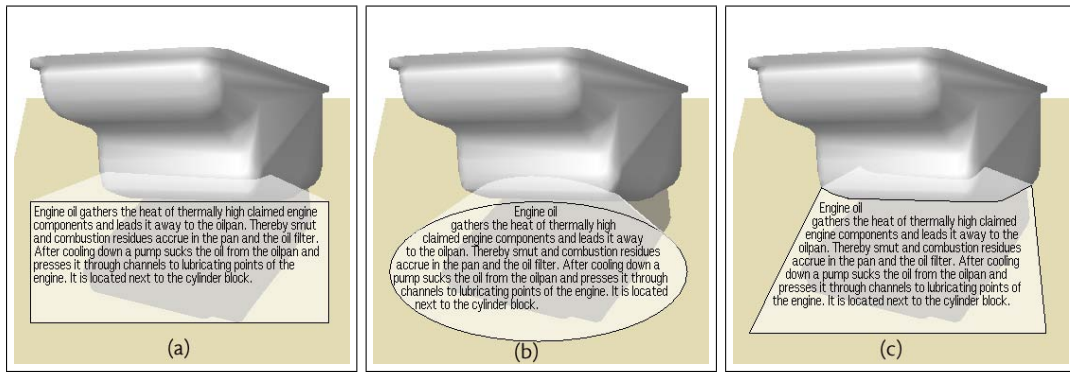


Figure 6: Figures (a) and (b) show outlined annotations whose rectangular bounding box is attached to an extreme vertex of the projected 3D object. In contrast, the annotation of Figure (c) is adapted to the corresponding object’s shape. Additionally, all displayed annotations are visually linked to their associated object by a transparent polygonal shape.

Rendering the annotation

The annotation is rendered as a 2D semi-transparent polygon in screen space. After the coordinates have been determined, the text is placed within the polygon by using a scanline algorithm that has been described by CHIGONA et al. [5].

Annotations are connected to their 3D objects through polygonal connector objects in 3D. These connector objects are rendered semi-transparently to make a clear link between the scene object, the connector polygonal object and the associated annotation (see ZHAI et al. [20]). While one part of these connector polygonal objects is visually attached to the annotation, the only vertex at the opposite side is located at the centroid of the 3D scene object (see, e.g., Figure 7(b)). This position has already been determined during the preprocessing step. In order to achieve visual connection between the remaining connector polygon vertices and the annotation, the required vertices of the annotation in 2D space have to be converted into 3D model space.

5.4 Interactively Accessing Annotations

After having chosen a suitable view of the 3D model, an object can be selected and its associated textual information can be explored by utilizing the distortion lens that is attached to it. The size of the appropriate annotation increases when the cursor is moved towards the center of the lens. During informal testing, we found that it is difficult to find the center position without seeing it. Hence, we modified the application in such a way that the selected object becomes transparent when it is entered by the cursor. Additionally, a sphere is rendered at the object’s centroid to indicate the center of the lens. As described in Section 5.2, a distortion results from moving the mouse cursor within the scope of the lens. Since the distortion can be used for both scaling the annotation and scrolling through the text, we had to find a solution for enabling the user to switch between those tasks. Figure 7 illustrates the procedure from selecting an object over scaling an annotation and keeping it in this state to scrolling through the text.

6. DISCUSSION AND FUTURE WORK

Figures 1 and 8 show typical use cases. Figure 1 illustrates the initial and later state during the exploration of the model

of a foot. Here a foot bone, that is initially not visible, appears when a probe is moved towards this area. Additional information can then be accessed as described in Section 5.4. Figure 8(a) shows an exploration of the heart. Here, the user is focused on the pulmonary arteries that appear completely after moving the 3D probe towards this area. Figure 8(b) shows a spatial exploration of an engine. Annotations remain attached to two different objects when further engine parts are explored. In all cases, the objects move back to their initial position as soon as the probe is moved to a different area.

The *Expanding Annotations* integrate more extensive textual explanations into the display of the 3D models. This of necessity requires a larger amount of screen space. Our approach handles this through: only displaying annotations when requested, only expanding to the degree requested and introducing scrolling when full text is requested and the necessary space is not available. Since it is possible to reveal other annotations while existing annotations remain in their state, annotations of adjacent model parts may overlap. This situation can also arise during other exploration techniques such as rotations. We are currently investigating possibilities of using partial transparencies and repositioning of annotations that are sensitive to the locations of other annotations.

While we have made use of informal discussions with potential users during the development of these techniques, we are designing a user study to more fully investigate the relative advantages and disadvantages of integrating text explanations with the display of 3D models.

7. CONCLUSION

In this paper, we have presented two methodologies for the interactive exploration of 3D models, a *3D Explosion Probe* and *Expanding Annotations*.

The 3D Explosion Probe can be moved through the model, providing access to the interior components of the model and enabling spatial exploration of the model. Using the probe to interact with the model affects all objects within the scope of the probe. This means that occluding objects can be moved smoothly to the side and then replaced. Objects do not have to be selected and moved individually.

The annotations expand smoothly on demand. They can




Figure 7: An expanding annotation. (a) Interest in an object is indicated with the mouse cursor. (b) When the cursor reaches the object, the object becomes transparent enough to reveal its centroid. (c) As the cursor is moved towards the centroid the annotation expands. (d) A second click highlights the object's shadow and now, moving the cursor will scroll the text. When the cursor leaves the object, the object's centroid is not needed and the object is no longer rendered as transparent.

be viewed at various stages of expansion as requested by cursor location. The user can control the shape of an annotation and whether this annotation will persist as the user continues with other explorations. Also, if the amount of available text requires still more space, the user can scroll through the text by moving the cursor within the annotation trigger area. We have combined those methodologies in an interactive environment for exploring 3D models.

8. SUPPORTING MATERIAL

A video that illustrates the research presented in this paper can be downloaded at <http://isgwww.cs.uni-magdeburg.de/~sonnet/ArbeitE/LensAPI>

9. ACKNOWLEDGMENTS

The authors would like to acknowledge support from *Heritage Canada*, *Intel* and *NSERC*. They would also like to thank the members of the *Interactions Lab* for helpful and insightful suggestions and comments. Finally, thanks go to Julia Schliebenow for her contributions to the 3D probe.

10. REFERENCES

- [1] B. Bell, S. Feiner, and T. Höllerer. View Management for Virtual and Augmented Reality. In *Proceedings of the 14th Annual ACM Symposium on User Interface Software and Technology*, pages 101–110. ACM Press, 2001.
- [2] M. S. T. Carpendale, D. J. Cowperthwaite, and F. D. Fracchia. Extending Distortion Viewing from 2D to 3D. *IEEE Computer Graphics and Applications, Special Issue on Information Visualization*, 17(4):42–51, July 1997.
- [3] M. S. T. Carpendale, D. J. Cowperthwaite, M. Tigges, A. Fall, and F. D. Fracchia. The Tardis: A Visual Exploration Environment for Landscape Dynamics. In *Proceedings of SPIE Conference on Visual Data Exploration and Analysis VI*, pages 110–119. SPIE, 1999.
- [4] M. S. T. Carpendale and C. Montagnese. A Framework for Unifying Presentation Space. In *Proceedings of the 14th Annual ACM Symposium on User Interface Software and Technology*, pages 61–70. ACM Press, 2001.

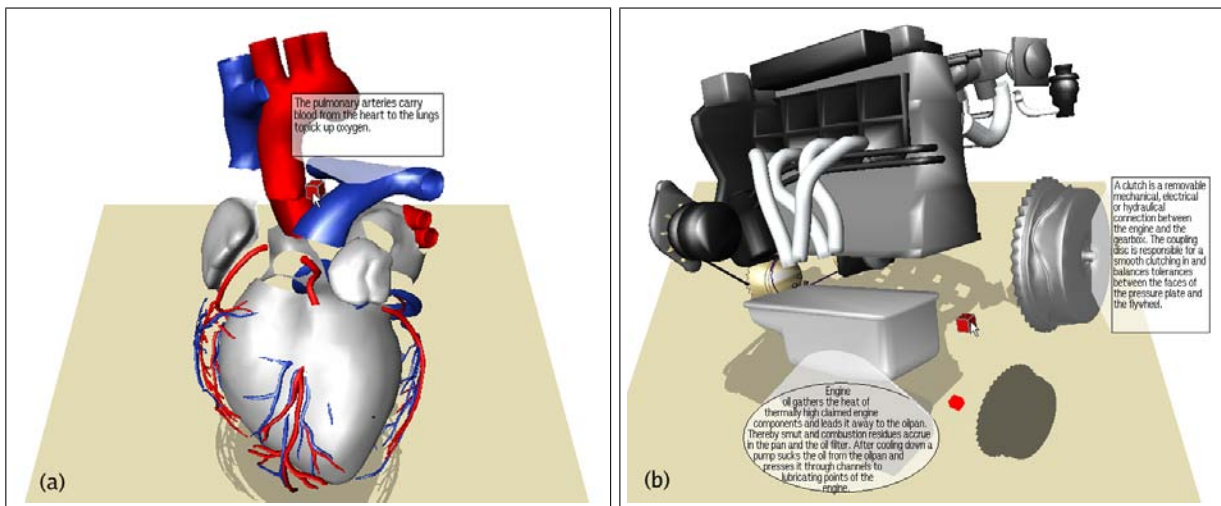


Figure 8: Two examples: (a) Exploring the model of a heart showing its pulmonary arteries that would be partially occluded in a non-exploded diagram. (b) Exploring the composition of an engine. In both cases, annotations are attached to selected objects.

- [5] W. Chigona, H. Sonnet, F. Ritter, and T. Strothotte. Shadows with a Message. In *Smart Graphics: Third International Symposium on Smart Graphics 2003*, pages 91–101. Springer Verlag, 2003.
- [6] O. Deussen, J. Hamel, A. Raab, S. Schlechtweg, and T. Strothotte. An Illustration Technique Using Hardware-Based Intersections and Skeletons. In *Proceedings of Graphics Interface '99*, pages 175–182. Canadian Information Processing Society, 1999.
- [7] J.-D. Fekete and C. Plaisant. Excentric Labeling: Dynamic Neighborhood Labeling for Data Visualization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 512–519. ACM Press, 1998.
- [8] G. W. Furnas. Generalized Fisheye Views. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 16–23. ACM Press, 1986.
- [9] T. Jung, M. D. Gross, and E. Y.-L. Do. Annotating and Sketching on 3D Web Models. In *Proceedings of the 7th International Conference on Intelligent User Interfaces*, pages 95–102. ACM Press, 2002.
- [10] Y. Kurzion and R. Yagel. Interactive Space Deformation with Hardware-Assisted Rendering. *IEEE Computer Graphics and Applications*, 17(5):66–77, 1997.
- [11] M. McGuffin, L. Tancau, and R. Balakrishnan. Using Deformations for Browsing Volumetric Data. In *Proceedings of IEEE Visualization 2003*, pages 401–408. IEEE Computer Society Press, 2003.
- [12] C. Niederauer, M. Houston, M. Agrawala, and G. Humphreys. Non-Invasive Interactive Visualization of Dynamic Architectural Environments. In *Proceedings of the 2003 Symposium on Interactive 3D Graphics*, pages 55–58. ACM Press, 2003.
- [13] A. Pommert, K. H. Höhne, B. Pflesser, E. Richter, M. Riemer, T. Schiemann, R. Schubert, U. Schumacher, and U. Tiede. Creating a High-Resolution Spatial/Symbolic Model of the Inner Organs Based on the Visible Human. *Medical Image Analysis*, 5(3):221–228, 2001.
- [14] B. Preim, A. Raab, and T. Strothotte. Coherent Zooming of Illustrations with 3D-Graphics and Text. In *Proceedings of Graphics Interface '97*, pages 105–113. Canadian Information Processing Society, 1997.
- [15] A. Raab and M. Rüger. 3D-Zoom: Interactive Visualisation of Structures and Relations in Complex Graphics. In *3D Image Analysis and Synthesis*, pages 125–132. infix-Verlag, 1996.
- [16] F. Ritter, H. Sonnet, K. Hartmann, and T. Strothotte. Illustrative Shadows: Integrating 3D and 2D Information Displays. In *Proceedings of 2003 International Conference on Intelligent User Interfaces*, pages 166–173. ACM Press, 2003.
- [17] M. Sarkar, S. S. Snibbe, O. J. Tversky, and S. P. Reiss. Stretching the Rubber Sheet: A Metaphor for Viewing Large Layouts on Small Screens. In *Proceedings of the 6th Annual ACM Symposium on User Interface Software and Technology*, pages 81–91. ACM Press, 1993.
- [18] R. Spence and M. Apperley. Database Navigation: An Office Environment for the Professional. *Behavior and Information Technology*, 1(1):43–54, 1982.
- [19] J. Viega, M. J. Conway, G. Williams, and R. Pausch. 3D Magic Lenses. In *Proceedings of the 9th Annual ACM Symposium on User Interface Software and Technology*, pages 51–58. ACM Press, 1996.
- [20] S. Zhai, W. Buxton, and P. Milgram. The Partial-Occlusion Effect: Utilizing Semitransparency in 3D Human-Computer Interaction. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 3(3):254–284, 1996.