Simplifying the Exploration of Volumetric Images: Development of a 3D User Interface for the Radiologist's Workplace

M. Teistler,¹ R. S. Breiman,² T. Lison,³ O. J. Bott,³ D. P. Pretschner,³ A. Aziz,⁴ and W. L. Nowinski¹

Volumetric imaging (computed tomography and magnetic resonance imaging) provides increased diagnostic detail but is associated with the problem of navigation through large amounts of data. In an attempt to overcome this problem, a novel 3D navigation tool has been designed and developed that is based on an alternative input device. A 3D mouse allows for simultaneous definition of position and orientation of orthogonal or oblique multiplanar reformatted images or slabs, which are presented within a virtual 3D scene together with the volume-rendered data set and additionally as 2D images. Slabs are visualized with maximum intensity projection, average intensity projection, or standard volume rendering technique. A prototype has been implemented based on PC technology that has been tested by several radiologists. It has shown to be easily understandable and usable after a very short learning phase. Our solution may help to fully exploit the diagnostic potential of volumetric imaging by allowing for a more efficient reading process compared to currently deployed solutions based on conventional mouse and keyboard.

KEY WORDS: Volumetric imaging, user interface, human-computer interaction, image overload, 3D visualization, 3D navigation, 3D input device, oblique reformation, slab reformation, volume rendering, maximum intensity projection, average intensity projection

INTRODUCTION

Volumetric imaging, ie, the utilization of data sets with isotropic or near isotropic resolution, has gained increasing importance in diagnostic radiology in recent years. Modern computed tomography (CT) and magnetic resonance imaging (MRI) scanners are capable of rapidly generating a large number of images per examination, allowing reconstruction of high spatial resolution volume data sets. With the advent of multidetector CT (64 and recently

256), the number of images per examination has drastically increased to hundreds and thousands. With examinations covering the whole body in relatively short scan times, multidimensional reformations of practically any area are possible. This availability of increased detail presents both an opportunity and a problem. On the one hand, it allows for more diagnostic options; on the other, the radiologist may no longer be able to efficiently view and interpret the resultant large number of images within given time constraints. For the current and future generations of imaging modalities, viewing the original source images may not be the best means of exploiting the full diagnostic potential of the data.¹⁻³

In an effort to maximize the diagnostic utility of modern cross-sectional imaging modalities, several visualization techniques for the display of volumetric image data are utilized, including multiplanar reformation (MPR), volume rendering

¹From the Biomedical Imaging Laboratory, Singapore BioImaging Consortium Agency for Science, Technology and Research (A*STAR), 30 Biopolis Street No. 07-01, The Matrix, 138671, Singapore.

²From the Department of Radiology, University of California, San Francisco, CA, USA.

³From the Institute for Medical Informatics, Technical University Carolo-Wilhelmina, Braunschweig, Germany.

⁴From the I-MED, Regional Imaging, Riverina Campus, Wagga Wagga, Australia.

Correspondence to: M. Teistler, u_m_i Informatik GmbH, Petritorwall 24, 38118, Braunschweig, Germany; tel: +65-647-88505; fax: +65-647-89049; email: teistler@virtusmed.info

Copyright © 2007 by Society for Imaging Informatics in Medicine

Online publication 27 March 2007

doi: 10.1007/s10278-007-9025-8

(VR), maximum intensity projection (MIP), minimum intensity projection (MinIP), and surface-shaded display (SSD).^{4,5} These techniques are well established and have been shown to be useful in promoting a better appreciation of the complex anatomic relationships often encountered in diagnostic imaging.⁶ Whereas much effort has been spent on the appearance of the resultant processed image itself and on optimizing performance issues, such as quality and speed of VR, there has been comparatively little emphasis on the user interface, ie, the question of how the radiologist or 3D technologist actually interacts with the image data. The 3D solutions that are currently deployed in radiology often add a level of complexity because of their cumbersome user interfaces, which may be an impediment to the widespread deployment of these potentially useful techniques. In clinical practice, most radiologists still prefer to look at large stacks of consecutive slices in a single plane for a vast majority of cases, at least in part because alternative methods are felt to be too complex and proficiency with 3D viewing is elusive to most users.

As a potential solution, the utilization of virtual reality techniques has been proposed for a relatively long time. Several tools have been developed that allow the user to explore volumetric data in a semi- or fully immersive environment. But currently these solutions are mainly used for research purposes and rarely in routine clinical settings. They often represent specialized solutions that do not directly address the interpretation of diagnostic radiology examinations, but rather focus on other areas such as simulation for preoperative planning (eg, see Kockro et al.⁷ and Reitinger et al.⁸). In addition, they tend to be relatively expensive because of the utilization of specialized hardware and the increased time, effort, and expense of installation, training, and maintenance. But most importantly, these techniques implement interaction and visualization metaphors fundamentally different from conventional solutions. As a result, they create a rather radical paradigm shift associated with a higher learning barrier than justified by the gains appreciated by most radiologists and 3D technologists.

The aim of this work was to develop a new 3D navigation tool for the exploration of volumetric image data that should be easy to learn and use, similar to current methods utilized for 2D image

manipulation and interpretation. The new technique should be of a generic nature so as to be applicable to a wide range of diagnostic imaging functions as opposed to solutions like virtual colonoscopy or virtual bronchoscopy that apply (semi-) automatic preprocessing on the image data to allow for an appropriate interactive visualization only within a specific, rather narrow context. In addition, to further increase acceptability, it should represent a low-cost solution with minimized hardware costs and minimal effort for installation and maintenance. Real-time interaction has to be provided, ie, low latency and high update rate, as a crucial prerequisite for an efficient interaction.

This paper describes the resulting computer system, which has been developed within the project "virtusMED" (<http://www.virtusmed.info>).⁹ The first prototype was presented at the annual meeting of the Radiological Society of North America (RSNA) in 2002 as the "Future Viewing Station."¹⁰ It was derived from an educational computer system, presented 1 year earlier at RSNA,¹¹ whose basic interaction metaphor was appreciated by radiologists and was supposed to have potential also for diagnostic purposes. The tool was iteratively improved and extended¹² and finally presented at the RSNA 2004 meeting, where it received very positive feedback as a solution addressing the problem of the "data explosion in radiology."¹³

METHODS AND MATERIALS

Basic Concept

The presented approach concentrates on interaction rather than visualization. For visualization, VR, MIP, and MPR are utilized, as they are already well-established techniques. As radiologists have experience and comfort with the review of a stack of 2D images, our system let the user still focus on 2D views based on slices or slabs, but provides a new way of 3D navigation to arbitrarily define their position and orientation within the image stack.

As for currently deployed solutions, the utilization of a mouse or trackball, which are limited to *two* degrees of freedom (DOF) for navigation, seems to be a major obstacle for exploring three-

dimensional data. It requires the user to switch between different modes of interaction, eg, between one for rotating and one for translating. For example, defining a double- or triple-oblique slice often requires multiple steps and becomes a nonintuitive, time-consuming task. Thus, an acceptable solution should utilize an input device with more than two DOF to provide maximum control of navigation through the volume image data set. But, as stated above, it should still be kept simple, cost-efficient, and resemble currently deployed solutions in radiology to maximize its acceptability and actual utility. For this reason a desktop-based approach was chosen over an immersive environment. The standard computer (2D) mouse was extended to 3D by a motion-tracking sensor to allow for direct three-dimensional interaction. Other than that, utilization of specialized hardware was avoided and off-the-shelf PC technology was used instead.

System Design

The user is provided with two views, usually on two monitors. A 3D view (Fig. 1, left) shows the full data set visualized with VR together with slices or slabs of selectable size and thickness and corresponding clip planes. For better orientation, a virtual model representing the input device and other objects, like a volume-related coordinate system, can be added and the edges of the slices/slabs can be color-coded. A 2D view shows the current slice or slab in plan view (Fig. 1, right). The slabs can be rendered using MIP (Fig. 2a–b), average intensity (ray sum) projection (Fig. 2c), or standard (alpha blending) VR technique (Fig. 2d). The VR of the full data set in the 3D view primarily serves as orientation and allows the user to better assess the spatial location of a slice or slab. In the following it will be referred to as the reference volume. To visualize multiple slices/

Fig 1. The 3D view (left) shows the volume data set (volume rendered, here: head and neck CT scan), reconstructed slice(s) with optional clipping planes(s) and some basic orientation markers (here: volume-related coordinate system). The 2D view (right) shows the current (active) slice in plan view.

slabs, the user can select the current one as a key image so it will remain in the 3D view until deleted.

Figure 3 illustrates the basic interaction metaphor. A 3D mouse is used to define position and orientation of the reference volume and of the slices or slabs. It works either as a conventional mouse (2D mode) or, when lifted from the desktop, provides simultaneously three translational and three rotational DOF, ie, six DOF (3D mode). In the latter case, the mouse movement, along the three spatial axes, and its rotation around these three axes are measured at the same time.

In the 3D mode, the mouse can be used for positioning a slice or slab in an either absolute or a relative manner. In the first case (Fig. 3a), the

mouse works like an ultrasound probe. The user examines a virtual volume situated on the desktop in a manner analogous to a real-time ultrasound examination. The reference volume is virtually located at a specific position in space and the (virtual) location of the current slice or slab image is directly and constantly linked to the (real) location of the mouse. In the second case (Fig. 3b), the position and orientation of the slice will be relatively adjusted according to the movement of the mouse only while the user is pressing a specific mouse button. After having started to hold the button down, the slice or slab's position and orientation will be changed according to the changes of the mouse's position and orientation. This allows the user to begin by holding the mouse

Fig 2. Exploring a head MRI data set with slabs. Maximum intensity projection in the 3D (a) and 2D view (b). 2D view of average intensity (ray sum) projection for the same slab (c). Standard (alpha blending) VR of a cube-like slab with corresponding clipping planes in the 3D view (d).

Fig 3. Illustration of the basic user interaction metaphor. A 3D mouse is used to change position and orientation of slices or slabs, either in an absolute (a) or relative manner (b, e), and of the reference volume together with the slices/slabs in a relative manner (c, d). Depending on the current task and the interaction mode, which can either be 3D (mouse lifted, a–c) or 2-dimensional (mouse on the desktop, d, e), different DOF are available.

in an arbitrary and convenient position when relocating the slice or slab because absolute position and orientation of the mouse do not matter. This relative positioning can also be applied to adjust position and orientation of the reference volume together with all existing slices and slabs (Fig. 3c). The center of rotation is always the center of the volume or the slice/slab, respectively.

In the 2-dimensional mode, apart from performing tasks related to the 2-dimensional graphical user interface (GUI) as usual, the user can rotate the reference volume together with all existing slices and slabs with conventional movements of the mouse (Fig. 3d). Moving the mouse forward or backward rotates the volume up or down. Moving it to the left or to the right rotates the volume to the left or to the right. In addition, the active slice/slab can be moved in a parallel way with only one DOF by moving the mouse forward and backward (Fig. 3e).

Allowing the user to absolutely position a slice or slab was motivated by the idea of adopting the well-known ultrasound paradigm to achieve an intuitive and easy-to-understand interaction. To support this (at least for the initial learning phase), the slice or slab can also be sector-shaped as are sliced images produced by a sector scanner, and the 2D view can also simulate an ultrasound-like view (see below). In contrast, the relative positioning provides a more flexible way of using the 3D mouse, thus potentially producing less fatigue and allowing for more accuracy in 3D navigation.

The combination of both a 2D and a 3D mode allows for a smooth integration of conventional with 3D interaction to let the user easily choose

the most appropriate interaction mode for the current task without the need of switching between different input devices. As an example, the user might want to use six DOF (lifting the mouse) for defining the orientation and initial position of a slice and afterward use only one DOF (putting the mouse back to the desktop) to scroll forward and backward between several slices parallel to the one first chosen.

Figure 4 gives an example of how the user is interacting with the system defining oblique slices in the absolute manner. An ultrasound probe model represents the input device that is used to define the location of sector-shaped slices. The 3D view also shows two planes representing two real markers on the desktop to help the beginner to understand the relationship between the virtual scene on the screen and the real-world interaction.

Several constraints can be applied to the six DOF interactions. The active slice/slab can be automatically adjusted to the nearest orthogonal orientation, either related to the volume coordinates to achieve axial, sagittal, or coronal orientations, or related to an already defined key image to create additional parallel or perpendicular slices. The user can also select different clipping modes. Clipping can be switched off, constantly at one specific side of the slice or slab, on both sides or always automatically on the side that faces the viewer (ie, the part of the volume between the slice or slab and the viewpoint is removed).

The 2D view can be computed in two different ways. In the “normal plain view” the slice or slab is shown in the same basic orientation as the slice in the 3D view (Fig. 1, right and Fig. 5c). In

Fig 4. Example of how to use the system (screenshots of the 3D view, absolute mode, sector-shaped slice). The input device lies on the table (1). The user lifts the device to explore the data set (2–3). The user marks a slice as key image and continues with another one (4–5). The user relocates the volume and the slices (zoom in) (6).

contrast, the “ultrasound-like view” is provided to carry the ultrasound metaphor further. The slice or slab is presented as if it would be a 2D ultrasound image, ie, with a locally fixed orientation so that the top of the 2D view is always showing the same edge of the slice/slab (which is marked by a small cone) (Fig. 5d). Thus, users with experience in ultrasound imaging can produce images similar to those they are used to. In addition, this approach has the advantage that the view is not dependent on the orientation of the reference volume in the 3D view. The first mode, on the other hand, makes sure that the current slice/slab is not shown in an orientation different from that one in the 3D view (to avoid, for instance, that the part of the slice that is shown on the left side within the 3D view is shown on the right side within the 2D view). A disadvantage of this approach is the dependence of the 2D view on the 3D view, which can cause

the 2D view to suddenly flip, when the reference volume is rotated or slices/slabs with an orientation nearly perpendicular to the 3D view direction are rotated (Fig. 5c, iii vs iv).

Movies can be downloaded that provide further examples of how to use the different features of the system (<http://www.virtusmed.info/movies/vmdiag>).

IMPLEMENTATION

The system is based on a standard PC with Microsoft Windows NT operating system (ie, Windows NT 4.0, 2000, or XP; Microsoft Corporation, Redmond, WA, USA). The software was implemented using Microsoft Visual C++ and the Microsoft Foundation Classes (MFC).¹⁴ For importing image data, the digital imaging and communication in medicine (DICOM) software

Fig 5. Different ways of visualizing a volume data set (here: abdominal CT scan) and an oblique slice. a 3D view without clipping. b 3D view with clip plane. c 2D view of the slice. d Ultrasound-like view of the slice. Each slice has a reference edge that is marked by a small cone. Here, triangles were added at the positions of these cones for better visibility.

library, DCMTK,¹⁵ has been integrated. For visualization, the graphics library Open Inventor¹⁶ in the version provided by Mercury Inc. is used, which provides a high-level application programming interface based on OpenGL.²¹ For visualizing the reference volume and the slabs, Open Inventor's extension VolumeViz¹⁷ is utilized, which implements VR using 3D textures as described in Wilson et al.¹⁸ The system deploys graphics cards with GeForce series 3, 4, or 6 graphics chips from NVIDIA Corporation, Santa

Clara, CA, USA that provide hardware acceleration of the OpenGL features used by Open Inventor, in particular, 3D textures. The 3D mouse is built by using a standard mouse and an electromagnetic motion-tracking sensor. Both Polhemus Fastrak¹⁹ and Ascension pciBIRD²⁰ have been used.

For conventional VR (alpha blending) and MIP, a simple transfer function for luminance and opacity is used, which allows for the definition of a threshold and adopts the typical windowing used

in radiological viewing. For the threshold, a partially quadratic function is used to compute the alpha values (Fig. 6). In our experience this has shown to provide well-defined boundaries but yet a smooth transition between visible and nonvisible voxels. The same luminance function is also used to determine the grayscale values of the 2D textures. Average intensity projection is implemented with Open Inventor's "sum intensity" VR. In this case, no threshold is applied, ie, no alpha value is used and the luminance values are divided by a value derived from the current slab thickness. Parameters like slice/slab width, depth and thickness, threshold, window center and width, window presets, clipping mode, and slice shape were made accessible through standard MFC-based GUI elements (sliders, radio buttons, and number input fields).

The 2D view is implemented by using a local virtual camera with fixed position relative to the center of the active slice/slab. For the ultrasound-like view, its orientation, defined by the view direction vector and the up direction vector, is also fixed. For the 2D view, these vectors are computed by taking the vectors belonging to the virtual camera that renders the 3D view and rotating them around the axis rotAxis by the angle rotAngle , which are computed as follows:

$$\begin{aligned} \text{rotAxis} &= \text{viewVector3DCamera} \times \text{sliceNormal} \\ \text{rotAngle} &= \cos^{-1} (\text{viewVector3DCamera} \cdot \text{sliceNormal}) \end{aligned}$$

where $\text{viewVector3DCamera}$ is the view vector of the camera in the 3D view, sliceNormal the normal of the active slice directing away from the 3D view point (so that $\text{viewVector3DCamera} \cdot \text{sliceNormal} > 0$), $\times$ the vector cross product, and $\cdot$ vector dot product.

RESULTS

The first prototype, presented at the RSNA 2002 meeting, followed the ultrasound paradigm, utilizing a transducer-like probe as input device (3D mode only) for creating sector-shaped slices as shown in Figure 4. To obtain an initial evaluation, the visitors testing the system were asked to fill out a simple questionnaire to judge the diagnostic usefulness of the presented approach and to optionally give comments. Forty-eight radiologists (49 nonradiologists, 9 not specified [n.s.]) filled out the questionnaire and judged the diagnostic usefulness of the system with an average of 4.46 (nonradiologists 4.45, n.s. 4.28) on a scale from 1 (not useful) to 5 (very useful). The adoption of the ultrasound paradigm was found promising compared to conventional solutions solely based on keyboard and 2D mouse. The sector shape was considered useful in the beginning to understand this new method of interaction with a volume data set and to appreciate spatial relations in the 3D view. After a short time using the system, however, this feature was

Fig 6. Illustration of the transfer function for standard VR (alpha blending), MIP, and oblique slices (only luminance).

considered unnecessary as it obscured portions of the anatomy at the periphery of the sector. At the RSNA 2004 meeting, the relative positioning mode, the combination of 2D and 3D interaction, and the definition of slabs were also presented. The visitors using the system were observed and informally interviewed. Relative positioning was felt to be useful because it allows for breaks when adjusting the location of a slice/slab and allows the user to choose any convenient position from which to start. Yet, absolute positioning seemed to be more intuitive because of the ultrasound analogy and because the spatial relationship between the input device and the slice/slab remains constant. The restriction of DOF when in 2D mode was considered to be a crucial feature to also have easy access to conventional one or two DOF interactions, which are supposed to be still the main way of interaction. The ability to create slabs was found to be an integral part of the system, allowing for evaluating a large portion of the data set at once by utilizing well-known visualization techniques, in particular, MIP. However, the radiologists surveyed had relatively little experience with oblique slabs and were not able to give more detailed feedback on the actual utility of creating them with an input device capable of six DOF.

In general, radiologists found the direct 3D navigation, ie, the ability to use all six DOF simultaneously to select reformatting planes or slabs and appropriate 3D views, easy to understand and use after a mostly very short learning phase. They appreciated the fact that they were provided with real-time interaction that gives them "full control" of the 3D display. In particular, this was found useful for the adjustment of an oblique slice/slab to follow an anatomical structure that does not conveniently lie in a sagittal, axial, or coronal plane (eg, tendons, optic nerves, petrous ridges, neuroforamina, and generally tortuous structures such as an ectatic aorta) and thus easily produce diagnostically relevant key images. It was also stated that the system would be very useful for the demonstration of imaging findings, with the potential to greatly improve the communication between radiologists and the referring physician.

The 2D view was mostly preferred to the ultrasound-like view by radiologists not experienced in sonography. The occasional flipping (as described above in the system design section) was initially confusing, but easily understood with some short

explanation. The system was tested with data sets of size up to $512 \times 512 \times 984$ voxels (984 slice images with 512×512 matrix). Using a PC with a Pentium IV CPU at 2 GHz, 2 GByte main memory, and a NVIDIA GeForce 6800 graphics card, latency and update rate were subjectively judged as acceptable. Also, no problems were encountered related to the accuracy of both of the chosen tracking devices.

DISCUSSION

By using six DOF for navigation, one relatively simple interaction metaphor can be used for a variety of tasks such as the determination of appropriate 3D viewing orientation, arbitrary slice/slab definition, and zooming. Using the latter one as example, the user normally has to define the zooming factor (one DOF) and the desired 2D region of interest (two DOF), utilizing a total of three DOF to perform this simple task. Because the normal mouse provides only two DOF, two consecutive steps are necessary. With the 3D mouse, the user literally grabs the slice and pulls it toward him or her.

The technique of manipulating 3D data sets with a combination of actions holding the device above the desktop and while resting on the desktop is natural and easy and likely to minimize user fatigue. The basic idea is to use 3D only when needed. This is not possible with immersive systems that often simulate 2D functions like pop-up menus within a 3D scene with a questionable usability.

The direct 3D navigation capability to arbitrarily define a (2D) slice or slab with a corresponding clip plane allows for an easy way of looking inside a volumetric image data set while making use of the experience and knowledge radiologists have in 2-dimensional image analysis. Volume rendering, on the other hand, has to utilize transfer functions ("color/opacity maps") or segmentation algorithms to allow for an insight. In these cases, a variety of parameters influence the visualization process as opposed to only window center and width for slice/slab viewing. This may increase the risk that the radiologist misinterprets the resulting image.

In general, the presented approach can help to fully exploit the diagnostic potential of high-resolution volumetric image data sets. It may promote new diagnostic approaches based on

oblique slices/slabs that otherwise would remain underutilized because they are too cumbersome and time-consuming to create. The ability to freely navigate through a 3D data set may result in many cases in an improved understanding by both radiologists and clinicians of the sometimes complex anatomic relationships between normal structures and pathology, potentially contributing to more accurate interpretation of imaging examinations and improved patient management. This may have applications in presurgical planning and in preprocedure simulation as well as intraoperative guidance. In this context, the focus is less on selecting a specific slice of interest as fast and accurate as possible, but rather on supporting the viewer's understanding of three-dimensional anatomic relationships by a dynamic real-time exploration that utilizes the (subconscious) link between hand movement and change in the viewer's perspective, providing a more lifelike display environment.

CONCLUSION

A novel solution for exploring volumetric images has been presented that is affordable, conceptually simple, and, according to the initial feedback from users, easily usable. It seems to fill an existing gap between conventional, keyboard- and mouse-based approaches and rather complex and expensive virtual reality systems by combining the advantages of both the 3D and the 2D worlds in an easily understandable way. It represents a viewing tool that appears to have applications in both existing and emerging volumetric imaging technologies, potentially making the process of interpretation and communication of results more efficient and accurate. However, further evaluation is needed to quantify and validate the clinical utility of this user interface approach to image manipulation. This includes the integration of this tool into existing picture archiving and communication systems (PACS) solutions to make it readily accessible to the radiologist during conventional interpretation of CT and MRI examinations.

ACKNOWLEDGEMENTS

The authors are grateful to all who tested the system and gave crucial feedback. Special thanks

goes to J. A. Brunberg, Department of Radiology, University of California Davis Medical Center, Sacramento, USA; J. Dormeier, Institute for Medical Informatics, Technical University Carolus-Wilhelmina, Braunschweig, Germany; N. Haramati, Department of Radiology, Albert Einstein College of Medicine and Montefiore Medical Center, New York, USA; D. Hirschorn, Department of Radiology, Massachusetts General Hospital, Boston, USA; Y. Rado, Department of Radiology, Heinrich-Heine-University, Düsseldorf, Germany; and C. R. Habermann, Center of Diagnostic Imaging and Intervention, Department of Diagnostic and Interventional Radiology, University Medical Center Hamburg-Eppendorf, Hamburg, Germany, for the fruitful discussions concerning the diagnostic utility of the "virtusMED" approach. Furthermore, M. Uetani, Department of Radiology, Nagasaki University School of Medicine, Nagasaki, Japan, and C. R. Habermann are gratefully acknowledged for providing useful medical images.

REFERENCES

1. Rubin GD: Data explosion: the challenge of multi-detector-row CT. *Eur J Radiol* 36:74–80, 2000
2. Andriole KP, Morin RL, Arenson RL, Carrino JA, Erickson BJ, Horii SC, Piraino DW, Reiner BI, Seibert JA, Siegel E: Addressing the coming radiology crisis—the society for computer applications in radiology transforming the radiological interpretation process (TRIP) initiative. *J Digit Imaging* 17(4):235–243, 2004
3. Rubin GD: 3-D Imaging with MDCT. *Eur J Radiol* 45(Suppl 1):37–41, 2003
4. Cody DD: AAPM/RSNA physics tutorial for residents: topics in CT. Image processing in CT. *Radiographics* 22:1255–1268, 2002
5. Addis KA, Hopper KD, Iyriboz TA, Liu Y, Wise SW, Kasales CJ, Blebea JS, Mauger DT: CT angiography: in vitro comparison of five reconstruction methods. *AJR Am J Roentgenol* 177:1171–1176, 2001
6. Haramati N: Interpretation strategies for large cross-sectional image data sets. In: Reiner BI, Siegel EL Eds. *SCAR University 2003, Educating Healthcare Professionals for Tomorrow's Technology*, The Society for Computer Applications in Radiology, VA, Great Falls, 2003, pp 169–172
7. Kockro RA, Serra L, Yeo TT, Chan C, Sitoh YY, Chua GG, Ng H, Lee E, Lee YH, Nowinski WL: Planning and simulation of neurosurgery in a virtual reality environment. *Neurosurgery* 46(1):118–137, 2000
8. Reitingen B, Bornik A, Beichel R, Werkgartner G, Sorantin E: Tools for augmented reality based liver resection planning. In: Galloway RL Jr Ed. *Medical Imaging 2004: Visualization, Image-guided Procedures, and Display*. Proceedings of the SPIE, volume 5367, 2004, pp 88–99

9. Teistler M, Bott O, Dormeier J, Pretschner DP: Virtual Tomography: a new approach to efficient human-computer interaction for medical imaging. In: Galloway RL Jr Ed. *Medical Imaging 2003: Visualization, Image-guided Procedures, and Display*. Proceedings of the SPIE, volume 5029, 2003, pp 512-519
10. Teistler M, Dormeier J, Dresing K, Franzen O, Habermann C, Bergmann J: The future viewing station: an intuitive and time-saving user interface beyond keyboard and mouse to improve CT and MRI based diagnosis (abstract). *Radiology Suppl* 225(P):764, 2002
11. Teistler M, Lison L, Dormeier J, Pretschner DP: Improving medical imaging understanding by means of virtual and augmented reality (abstract). *Radiology Suppl* 221(P):731, 2001
12. Jabs M, Saboor S, Lison T, Teistler M, Pretschner DP: How to read CT and MRI images with novel 3D techniques—managing exploration paths to improve the diagnostic process (abstract). In: *RSNA '03 Scientific Assembly and Annual Meeting Program*, Radiological Society of North America, Oak Brook, IL, 2003, p 805
13. Teistler M, Nowinski WL, Rado Y, Breiman RS, Bott OJ, Lison T: Data explosion in radiology: solving the problem on the user interface side (abstract). In: *RSNA '04 Scientific Assembly and Annual Meeting Program*, Radiological Society of North America, Oak Brook, IL, 2004, p 831
14. Jones RM: *Introduction to MFC Programming with Visual C++*, Microsoft Technologies Series. NJ: Prentice Hall, 1999
15. DCMTK-DICOM: DCMTK-DICOM-Toolkit, Kuratorium OFFIS e. V. Oldenburg, Germany: DCMTK-DICOM. Available at <http://dicom.offis.de/dcmtdk.php.de>
16. Wernecke J: *The Inventor Mentor, Programming Object-oriented 3D Graphics with Open Inventor*, Release 2. Ontario, Canada: Addison Wesley Longman, 1994
17. TGS Inc: *Open Inventor VolumeViz*. San Diego, CA: TGS Inc., 2002. Available at <http://www.tgs.com/support/datasheet/VolumeViz.pdf>
18. Wilson O, Gelder AV, Wilhelms J: *Direct Volume Rendering Via 3D Textures*. Technical Report UCSC-CRL-94-19. Santa Cruz: University of California, 1994
19. Polhemus Inc: *Fastrak*. Colchester, VT: Polhemus Inc., 2004. Available at http://www.polhemus.com/?page=Motion_Fastrak
20. Ascension Technology Corporation: *pciBIRD*. Burlington, VT: Ascension Technology Corporation, 2005. Available at <http://www.ascension-tech.com/products/pcibird.php>
21. Neider J, Woo M: *The Official Guide to Learning OpenGL, Version 1.2*. OpenGL Architecture Review Board, Addison-Wesley Longman, Amsterdam, 1999