J Digit Imaging (2012) 25:792–801
DOI 10.1007/s10278-012-9461-y

# Interactive Visualization and Analysis of Multimodal Datasets for Surgical Applications

**Can Kirmizibayrak · Yeny Yim · Mike Wakid · James Hahn**

Published online: 14 February 2012


**Abstract** Surgeons use information from multiple sources when making surgical decisions. These include volumetric datasets (such as CT, PET, MRI, and their variants), 2D datasets (such as endoscopic videos), and vector-valued datasets (such as computer simulations). Presenting all the information to the user in an effective manner is a challenging problem. In this paper, we present a visualization approach that displays the information from various sources in a single coherent view. The system allows the user to explore and manipulate volumetric datasets, display analysis of dataset values in local regions, combine 2D and 3D imaging modalities and display results of vector-based computer simulations. Several interaction methods are discussed: in addition to traditional interfaces including mouse and trackers, gesture-based natural interaction methods are shown to control these visualizations with real-time performance. An example of a medical application (medialization laryngoplasty) is presented to demonstrate how the combination of different modalities can be used in a surgical setting with our approach.

C. Kirmizibayrak · Y. Yim · M. Wakid · J. Hahn
Department of Computer Science,
The George Washington University,
Washington, DC, USA

C. Kirmizibayrak · Y. Yim · M. Wakid · J. Hahn
Institute for Biomedical Engineering,
The George Washington University,
Washington, DC, USA

C. Kirmizibayrak (✉)
Department of Radiation Oncology, Stanford School of Medicine,
875 Blake Wilbur Drive,
Stanford, CA 94305, USA
e-mail: kirmizi@gwmail.gwu.edu

**Keywords** Volume visualization · Human–computer interaction · Volume rendering · Image-guided surgery

## Introduction

Volumetric medical datasets are ubiquitous, but visualizing desired information, filtering out irrelevant information, and understanding the spatial relationships between them are challenging tasks, especially with multiple information sources [1]. Examples include data sources such as CT, MRI, PET, and SPECT, as well as color information from endoscopic cameras that are not present in volumetric modalities, and computer simulations producing 3D vector fields such as blood or air flow [2]. Inter- and intra-modality occlusion of data values complicates the visualization of the 3D information on a 2D screen. Another important feature surgical visualization approaches need to have is helping the surgeon to construct the cognitive mapping between the virtual world and the real world. The physician has to deduce the correspondence between these different spaces, which is usually helped by his domain knowledge. This manual mental registration might result in large variability and errors. For instance, incorrect judgment of surgical target location might lead to suboptimal surgical results or unnecessary trauma to neighboring anatomical structures. Even though volumetric visualization approaches aim to provide accurate 3D information to the surgeon, spatial relationships between these datasets can be difficult to comprehend. For example, the perception of one volumetric object passing behind another is based on mental models of object shape and movement, rather than direct volumetric perception [3]. Therefore, interactive volume visualization systems that provide explicit spatial control can be useful to infer these relationships.

Our visualization approach aims to overcome these problems, with the goal of integrating multiple data modalities in a single view, to help with establishing the cognitive mapping among them. The system gives the user a detailed rendering of the "interesting" parts of the data (the focus region), while displaying the rest (the context region) with less emphasis to provide the contextual information for the focus. However, deciding what is "interesting" is a difficult problem, as this can change for each application domain and user. In our case, the interesting parts can be a subset of one or more of the datasets. We give the user interactive control over the location of focus region (which is a sub-volume with either static or user-defined shape) and the rendering parameters for this region (e.g., applied transfer function). We define these focus regions as "lens filters". A lens filter can be assigned a dataset and corresponding rendering parameters (e.g., transfer function, transparency). Furthermore, as will be described in the following sections, analysis operations (e.g., calculating local histograms) and non-traditional medical datasets (e.g., vector-based, color information) can be assigned to different lens filters. Multiple datasets and rendering styles can be combined in one or more different lens filters, since our approach allows selection of multiple arbitrarily shaped volumetric regions. This approach has two main advantages. First, the user can change the rendering style for a specific region of interest and using a combination of these styles can create a desired visualization effect. Second, by interactively exploring and manipulating the datasets, information, and complex relationships between datasets that are not apparent can be discovered, while maintaining the understanding about spatial relationships.

The purpose of this paper is to present a conceptual overview of our visualization approaches and to describe how they can improve on the current medical practice. We have focused on surgical applications in our research because the contextual information is very important to help establish the connection between the visualization and the patient, since the main goal of the visualization is to reach a certain surgical target on the patient. Therefore, the surgeon has to make the connection between what he sees on the screen and the patient (i.e., mental registration). Furthermore, color information that is normally not available in medical modalities is fused into the visualization to help with this cognitive mapping. We envision that focus+context visualization methods can provide a viable alternative to traditional 2D-slice-based visualization methods because they can produce 3D visualizations, which require an easier cognitive mapping to the patient since they are more similar to what the surgeon sees in the operating room. At the same time, the lens filters can be used to remove occlusion in local regions, which is one of the most important obstacles in volume visualization, and reveal desired internal anatomical information.

Image-guided surgery systems have been applied to various surgical procedures [4, 5]. Focus+context visualization is an information visualization technique that has been applied to medical visualization by several researchers, examples include ClearView [6] and Svakhine et al.'s illustrative visualization methods [7]. A subset of focus+context visualization called magic lens uses static shaped spatial filters to define the focus region [8] and has been applied to volume visualization [9, 10]. Our application extends this method to handle real-time volume selection operations using multiple medical datasets (volumetric, image, and vector-based). This approach is similar to a volume editing method, such as VolumeShop [11] and Burger et al.'s direct volume editing approach [12]. In this paper, we build on our previous work in improved performance and removing limitations of previous volume editing methods such as limited brush sizes [13]. First, we propose an extension to the lens filter approach by introducing analysis lens filters, which can perform computations in user-defined local regions. We also propose to add datasets such as vector-based 3D datasets and color information from 2D images to be combined with volumetric datasets. These datasets are traditionally displayed separately from 3D renderings in medical practice, requiring the physician to mentally combine these multiple sources of information. We believe that combination of these multiple datasets can be an improvement to the current medical practice, and visualizing all available datasets in a single view can improve the understanding of spatial relationships between multiple datasets. The following sections will describe conceptual details about our magic lens visualization and volume editing approach, introduce our data analysis and vector-based dataset lens filters and their implementation, discuss registration and fusion of 2D images with 3D datasets, and introduce interaction methods we propose to use in our applications.

## Materials and Methods

### Magic Lens Visualization

Magic lenses were first defined as spatial filters to "modify the visual appearance of objects, enhance data of interest, or suppress distracting information" [8]. Our application uses multiple co-registered datasets to enable data exploration and improve understanding of spatial relationships between different datasets. Figure 1 demonstrates the use of this visualization method where internal structures of volumetric datasets are visualized together with the surrounding context of skin and soft tissue, with the user exploring the dataset by changing the location of a sub-volume. Since the context rendering is similar to what the surgeon would see in the operating room, mental registration is improved. Furthermore, since the lens region has a fixed 3D shape and size, interactively changing its location can improve spatial understanding of rendered anatomical structures.

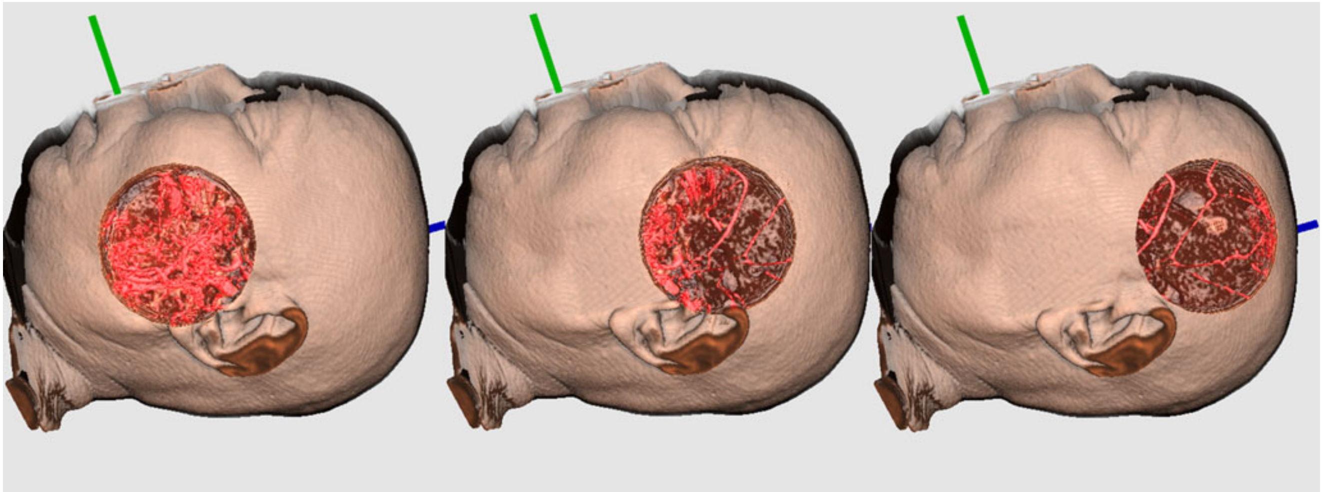

**Fig. 1** Example of use of magic lens visualization for data exploration by moving the lens left to right. The lens region uses a transfer function that shows vascular structures, while the rest of the dataset shows skin and soft tissue to provide context

## Interactive Region Selection and Multimodal Visualization

Magic lens visualization can be beneficial for data exploration; however, the lens shape is usually static. A useful extension to this paradigm is enabling real-time region selection. This way, the lens volume is used as a volumetric brush to change the rendering styles in arbitrary-shaped user-selected regions. Different lenses can be assigned different dataset/transfer function pairs, and the users can select multiple arbitrary shaped regions to see the spatial relationships contained in different datasets. We can use this approach to visualize multiple datasets in a single view. Figure 2 is an example of this: 3D renderings of co-registered MRI, CT, and PET scans are combined in a single view to combine information from these multiple information sources in a cohesive view. The lens region has been used as a brush to perform an accumulated selection action. The advantage of this approach is that the lens area can be considered as a "preview" of this editing action. While exploring the dataset with the magic lens, the user can decide to save the rendering parameters to the currently selected region, resulting in arbitrarily shaped selection regions. This way, the exploration, and editing tasks are conceptually combined, and the user can interactively toggle between volume exploration and volume selection (or editing). It should be noted that an ICP-based registration method [14] was used for intra-modality registration in our examples, but the exact nature of this registration is not important for our visualization method.

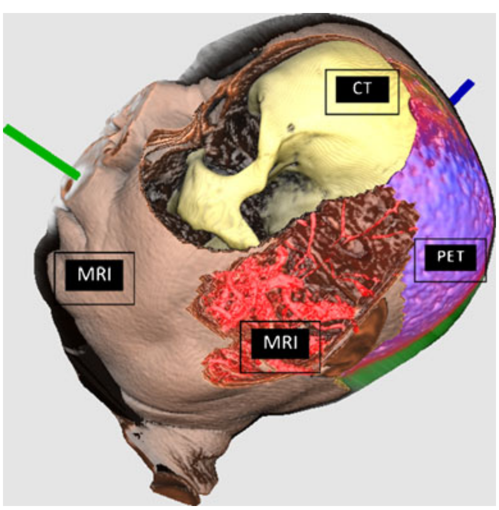


**Fig. 2** Example of multimodal visualization in a single view, the user has selected arbitrary shaped volumetric regions to combine CT, MRI, and PET datasets

## Local Image Analysis

Data exploration using volumetric lens filters enables the user to change the visualization locally. A useful extension to this is to give user data analysis tools applied to the same local region. Data features that are difficult to comprehend visually can be visualized with additional data plots, for instance by displaying histograms and scatterplots of selected local regions. The lens filter for visualization and analysis can be same or different, e.g., the user may want to use a lens filter to see the soft tissue from an MRI, while seeing a histogram of the CT dataset in the same region. Figure 3 shows an example analysis lens filter: a local histogram.

Since the location of the lens filter may change at any time, the analysis lens filters have to be recalculated in every frame. Even in the case of histogram, which is algorithmically very simple, this introduces a big computational burden, as every voxel contained in the lens filter has to be considered, and this computation has to be repeated every frame while keeping interactive frame rates. We have developed an implementation

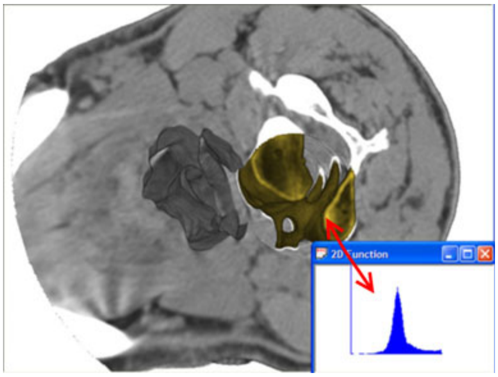

**Fig. 3** A data analysis lens filter, showing the histogram of the currently selected lens region

that takes advantage of modern graphics processing unit (GPU) architectures (specifically vertex texture fetch architecture [15]), which enables interactive local computation for analysis filters, inspired by the scattering technique to calculate histograms [16]. From a high level, this approach uses the fact that drawing simple point primitives are fast operations that can be processed in parallel by the help of GPUs. Figure 4 illustrates this approach using a 2D image for simplicity. To calculate a histogram, a point is drawn on a 1D image for each voxel in a 3D dataset. The location of where to draw each point is based on the voxel value. If more voxels contain a certain value (e.g., value "1" in Fig. 4), more points are drawn at that location in the histogram image. After all the points are drawn on the histogram image, a simple blending operation accumulates multiple points drawn, and creates the resulting histogram image. For a 2D scatterplot, a similar process would be carried out using a 2D output image rather than a 1D histogram.

## Visualization of Vector-Based Datasets

In a growing number of medical applications, vector field-based datasets are being used along with scalar ones, such as acquired datasets showing change (e.g., Doppler, 2D/3D contrast-enhanced ultrasound) as well as computer simulations (blood or air flow using computational fluid dynamics (CFD)). Displaying this information along with the anatomy would greatly improve the understanding between vector fields and their anatomical context.

As an example for such an integration, we introduce a vector field visualization lens filter (Fig. 5). We used a flow volume dataset produced by a CFD simulation. The CFD data was acquired using the same volumetric geometry; therefore, registration was not an issue. The flow was visualized using particle advection. The lens filter can be used to change the color or opacity of particles inside and outside the lens region. For instance, color can be changed with respect to velocity, vorticity, or pressure, and the opacity can be selected to reflect the magnitude of these values. Any of these options can be combined in the same or using multiple lens filters. For instance, a user can select one filter to highlight high-velocity data points while a second filter can make low-pressure particles transparent. This local selection can be used to reduce occlusion caused by particles in less important locations and highlight the desired information.

In particle advection-based flow visualization techniques, it is necessary to keep "injecting" particles into the flow stream to maintain visualization continuity. However, introducing too many particles can result in cluttering and data occlusion. We implemented an interaction scheme to interactively change the injection position. By changing the location of injection, the user can inject particles in a local region to help improve the local understanding of flow features in the desired region. An example for efficient use of this interaction technique is virtual endoscopy. By injecting flow at the current camera position, the particles can be injected to the current viewport, reducing cluttering, and occlusion due to unnecessary particles.

For better integration to our raycasting shader, we chose to use a particle-based visualization method. We used a similar approach that was briefly outlined by Li et al. [17]. The vertex shader was used with two passes to render the flow particles. The particles are initialized at positions interactively selected

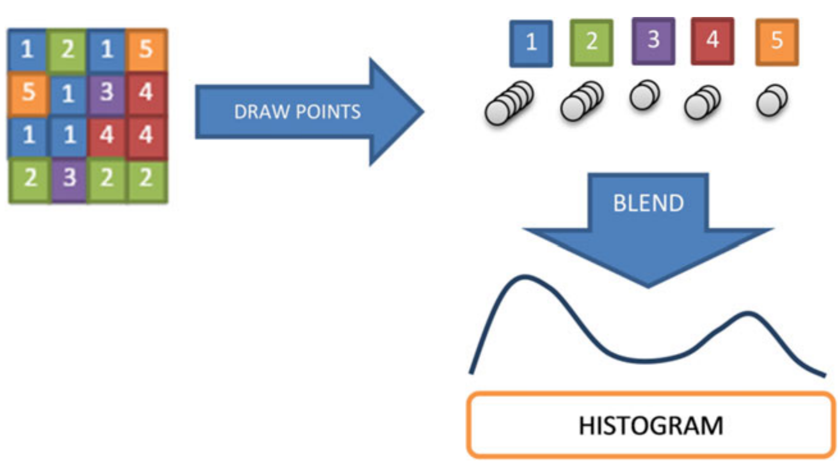


**Fig. 4** Diagram of histogram creation, illustrated on a 2D image for simplicity. Point primitives are drawn on the histogram image based on their intensity values, and simple blending is applied to get the histogram image

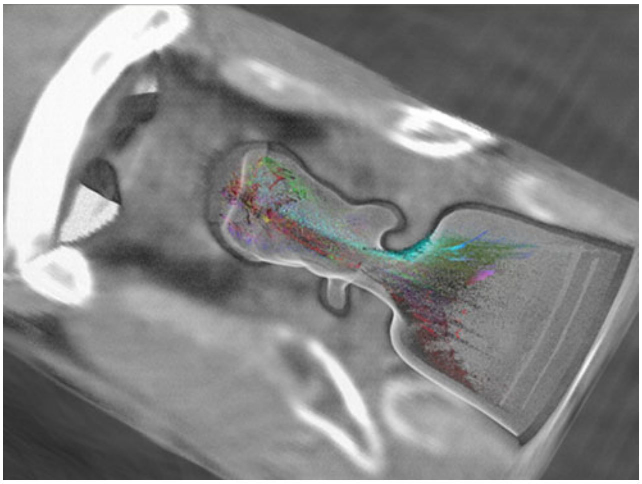

**Fig. 5** Example of vector field lens filter, where particles are colored according to their velocities. Anatomical structure of the airway is shown transparently to help with the understanding of spatial relationships

as explained above inside the volume, and the position for each vertex is stored in a texture. In the first pass, by fetching velocity values for the corresponding positions, particles are advected to their new positions. The second pass is done to render the particles, and appropriate opacity/color values are assigned based on the selected lens filter style and current lens filter location. To maintain interactive frame rates, total number of particles was kept constant, and particles are "recycled" back into the flow when they (1) exit the flow volume or (2) are "stuck" with near-zero velocity positions. This approach produces real-time frame rates for flow visualization and can be easily combined with traditional raycasting-based volume rendering techniques.

## 3D Image Fusion

As described in the Introduction, one of our main motivations of using a focus+context visualization-based visualization method was to help with the mental registration process in the operating room. Color and texture information is crucial to establish these kinds of connections, but are not present in traditional medical datasets such as CTs and MRIs. In a growing number of applications, still images or video frames from external (e.g., maxillofacial surgeries) or internal (e.g., endoscopic surgeries) are present, but displayed separately from the volume visualizations, complicating the mental registration process. We propose to use 3D image fusion to merge the color and texture information from video frames with the topographic information from the 3D CT or MRI. By doing so, the spatial relationship of information present in both imaging modalities can be visualized in a single coherent view.

Using the registration method that will be described in the next section, a corresponding rendered virtual view is created for each video frame. Since these views are registered, we should have a correspondence between every virtual point in the real camera image, from which the color information will be taken. Then, the color information is projected onto the corresponding surface in the rendered view, and this correspondence is stored in texture coordinates of the polygonal meshes that are created from the volumetric dataset in a preprocessing step. The video image is then texture mapped onto the surface where it remains there permanently. Weighted blending is applied when multiple video frames overlap on a surface. Weights are assigned to each image for each polygon in the polygonal mesh based on the polygon's distance and angle with respect to the camera during image projection.

After the fusion process, the user can select arbitrary viewpoints to visualize the video in relation to the patient's 3D volume, which is an improvement compared to traditional augmented reality-based approaches. From the selected viewpoint, a modified version of traditional raycasting-based rendering is performed to combine the color texture mapped on the surface mesh with volume samples to produce a final volume-rendered view of the fused data. Figure 6 shows an example of this, where a volume rendering is created from a single endoscopic video frame fused with a 3D CT data set. An extension to our lens filters approach is using lens filters to control the transparency of the texture-mapped images locally, which can be used to display internal structures in desired areas while keeping color information in the rest of the visualization to give context.

## Registration of 2D and 3D Modalities

To be able to fuse 2D and 3D images, we have to produce virtual views that are similar to our video frames (e.g., Fig. 6(a) and (b)). For this 2D-3D registration, we applied a previously proposed method [18].

The registration is performed by matching the viewpoints of the two cameras; one is a real camera, which captures the 2D

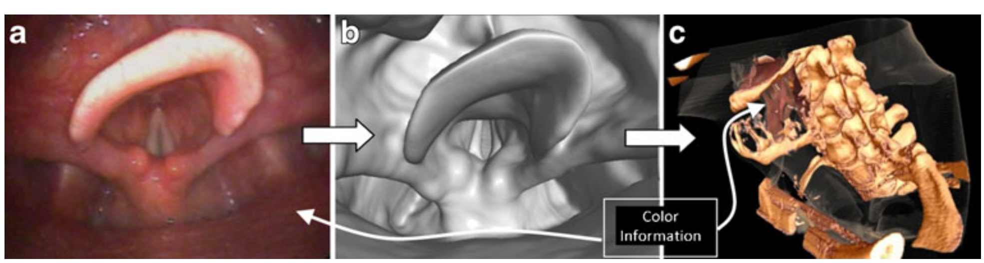


**Fig. 6** Volume rendering (**c**) of a single endoscopic video frame (**a**) fused with a 3D CT data set (**b**). Note that after fusion is performed the volume rendering can be viewed from any arbitrary viewpoint

images, and the other is the virtual camera, which is used to render the volume from 3D CT scans. Optimal viewpoint is determined by finding the viewpoint with the maximum similarity measure between the 2D image and the CT rendered image. We used normalized mutual information (NMI) as a similarity measure [19]. To find an optimal viewpoint with the largest NMI, we adapted downhill simplex optimization method. Although NMI is a robust measure for multimodality registration, the registration based on NMI can be unstable when it is applied to the 2D images and CT rendered image that have quite different surface illumination. For a robust registration regardless of the difference of the two images, the NMI is calculated in the gradient images as well as in the original intensity images. The viewpoint is optimized using the weighted sum of the two NMI values. To emphasize the measure calculated from the region of interest, the high gradient pixels in two images are found and a larger weight is assigned to the region of interest for the calculation of the NMI.

### Interaction Methods

The final component in our visualization approach is defining the interaction methods to be used by the users to control the lens location. The interaction with the system can be done in a variety of ways: using the mouse (where the projected cursor position is used as the current lens position), using optical or electromagnetic trackers (where the position and orientation of tracked sensors/markers can be used in 3D to define the lens transformation), or using the extracted hand positions of the user with a depth camera (e.g., Microsoft Kinect). Figure 7 shows examples of optical and electromagnetic tracking, and gesture-based interaction using the user's hand position. Each of these methods has advantages and disadvantages. The mouse is familiar and cheap, but translating the 2D motion of the mouse to 6 degrees-of-freedom of the lens transformation is difficult. Trackers are generally used in surgical settings, but cost, line-of-sight requirement, and calibration (for optical tracking) and possible interference (for electromagnetic tracking) can cause problems. Furthermore, all these tactile interaction methods cause sterilization concerns when used in the operating room environment [20]. Gesture-based interaction eliminates the sterilization requirement because the camera can be placed in the non-sterile working area and no additional equipment is necessary. Furthermore, intuitive interaction methods can be developed since people are naturally used to interact with their environment using gestures. For instance, our recent studies have shown that novice users can perform volume rotation tasks more effectively using a gesture-based interface compared to the mouse [21]. However, accuracy of such methods has to be studied further: we will provide our results comparing the accuracy of using the mouse and a gesture-based interface using the magic lens visualization in the results section.

## Results

The proposed methods have been implemented to provide interactive frame rates. A Dell Precision 690 system with 4GBs of RAM with an NVIDIA Quadro FX 4600 graphics card with 768MBs of graphics memory was used to test the rendering speeds. The introduction of the lens rendering causes no noticeable performance drop, since the only pre-rendering

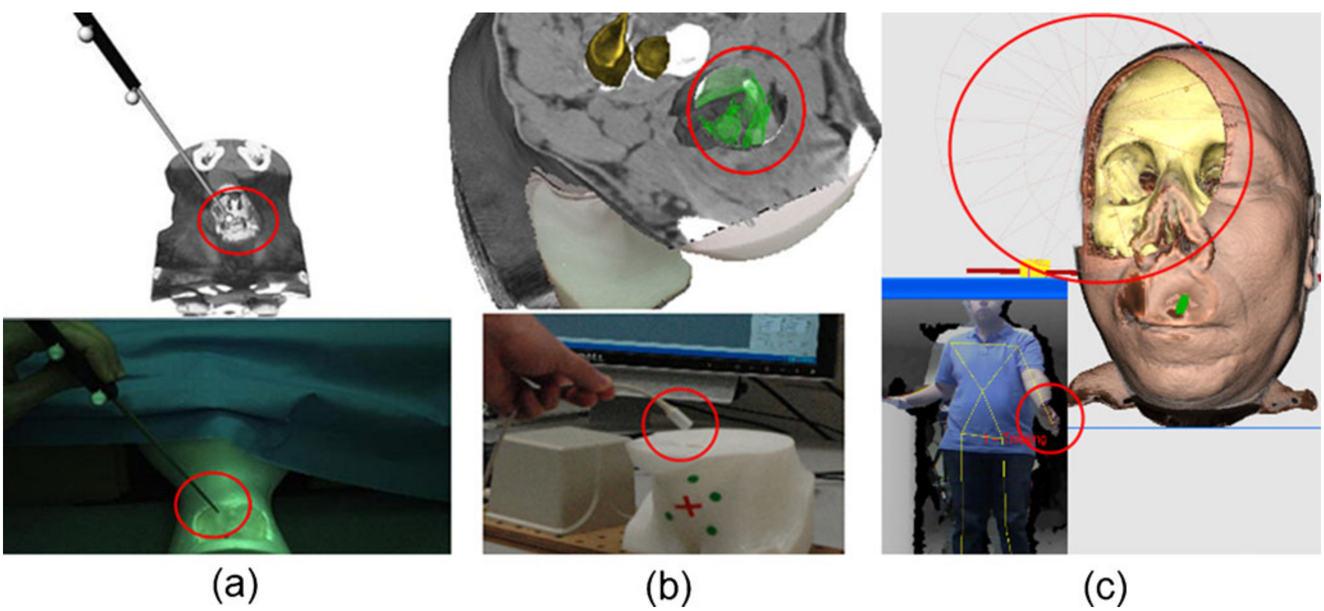


**Fig. 7** Interfaces used for interaction, (**a**) optical tracking, (**b**) electromagnetic tracking, (**c**) gestures (extracted hand location)

step is rendering the polygonal lens shape (about 0.5 ms, with 30 frames per second overall performance). Volume editing performance was achieved about 92% of volume rendering frame rates (around 26 frames per second). Similarly, GPU implementation methods enable image fusion operations, analysis, and vector-valued lens filters to be applied interactively.

We have evaluated the registration accuracy by measuring the sensitivity to the initial viewpoint of the CT rendered image. The registration was initialized by fixing the real camera at a true viewpoint which was calculated by manually matching the CT-rendered image to the endoscopic image as close as possible. Starting from multiple initial viewpoints of the virtual camera that were within a vicinity of the true viewpoint, an optimization process found the optimal viewpoint. The sensitivity to the initial viewpoint was measured by calculating the error between the true viewpoint and the registration result and by finding the range of the initial viewpoints that result in the acceptable registration. If the error is low for a wider range of the initial viewpoints, the registration can be considered as more accurate and acceptable. We assumed that the registration is acceptable when the difference of the view positions is less than 6 mm. Previous research suggested that 5 mm of positional and 5° of rotational error is an acceptable threshold [22], therefore we tested if our methods can produce results in this acceptable range. Our tests indeed showed that the registration was acceptable for the initial viewpoints translated by (−8, 8 mm), (−10, 10 mm), (−7, 6 mm) along *x*-, *y*-, and *z*-axis and rotated by (−20°, 11°), (−7°, 14°), (−20°, 20°) along *x*-, *y*-, and *z*-axis.

We have conducted a user study comparing a slice-based 2D visualization system using the mouse with the magic lens visualization using gestures extracted via Microsoft Kinect camera. 15 volunteers participated in the study (average age 29.4, all college-educated adults that consented to an IRB-approved experiment protocol). We used the OpenNI framework [23] to extract the hand position for the experiment. The experiment consisted of finding artificially inserted targets placed inside a volumetric dataset. The mouse interface used a slider control that changed the location of the slice (shown on the right side of the screen, left side showing the volume rendering showing the overall shape and a placeholder for the slice). The lens visualization showed the same dataset in 3D, where the lens volume revealed the targets in the currently selected location. Example screens can be seen in Fig. 8. Each trial had a randomly selected target and nine additional distractors (same shape as the target with smaller size), and the performance was compared in terms of time and accuracy. The magic lens interface was faster on average (mean, 6.6 s) compared to mouse (8.2 s), but had a higher variance (2.8 and 2.0, respectively). In terms of accuracy, mouse outperformed the gesture-based interface (5.3 vs 9.9 error units). Data distribution plots can be seen in Fig. 9. Even though our relatively small sample size might have prevented us from achieving more definitive results on the effectiveness of magic lens and gesture-based interfaces, the analysis of the data provided us with valuable and promising results. The experiments reinforced our assumptions that focus+context visualization can be used for fast exploration of volumetric datasets. However, for successful use of gesture-based interfaces in volume visualization tasks, further research into methods to improve accuracy is necessary.

## Discussion

### Example Application: Medialization Laryngoplasty

As a test bed to demonstrate the effectiveness of our approach, we used pre-surgery planning and image guidance for medialization laryngoplasty [24, 25]. This surgical procedure aims to correct vocal fold deformities by implanting a uniquely configured structural support in the thyroid cartilage. The implant shape and location is very critical, which makes the revision

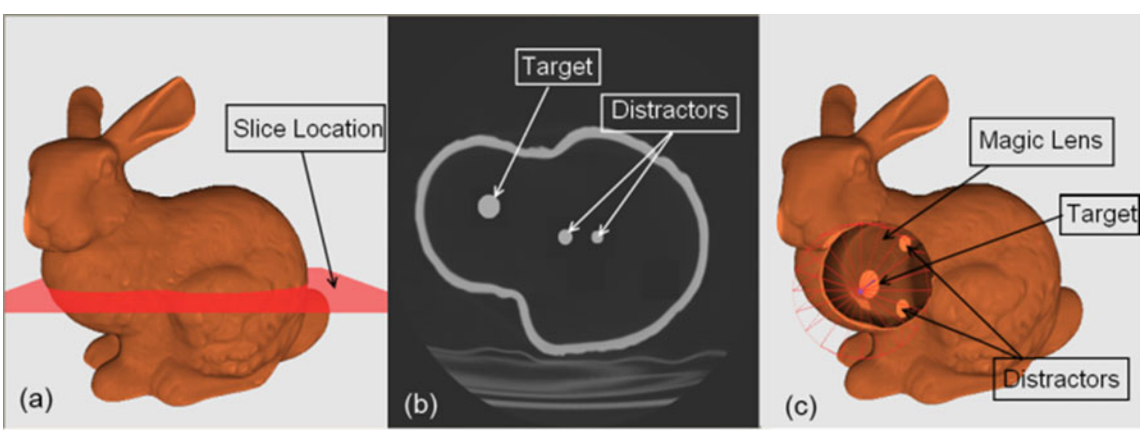


**Fig. 8** Sample screens used for experiments

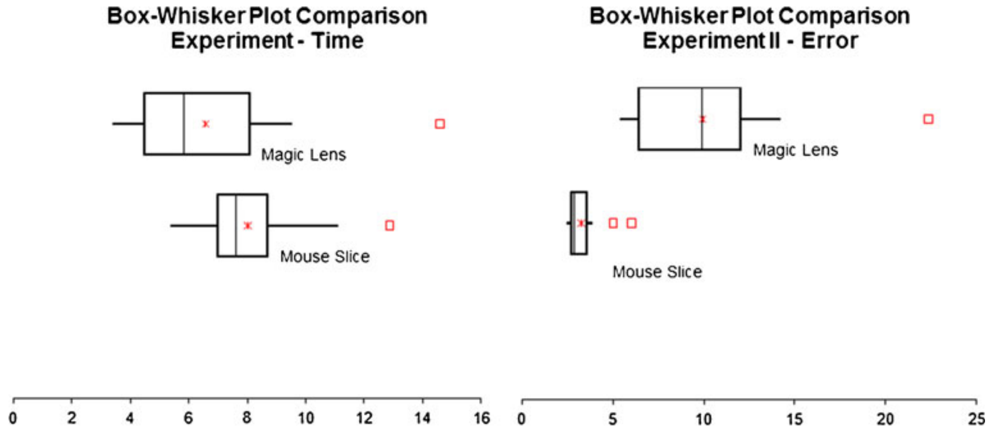


**Fig. 9** Data distribution for the task completion time and accuracy comparing magic lens and mouse-controlled 2D-slice-based visualizations

rate for this surgery as high as 24% even for experienced surgeons. In current medical practice, a preoperative CT scan and laryngoscope (a non-rigid endoscope inserted through the nasal cavity) evaluation is done prior to the surgery; and an implant shape and location is planned. During the surgery, the surgeon uses this information, along with the revealed anatomy and intra-operative laryngoscope to place the implant for optimal vocal fold correction. Vocal fold deformities causing problems for phonation are analyzed using fluid flow generated by patient-specific CFD simulations [2, 26]. Our choice of this procedure was motivated by the number and type of data modalities used in the decision making: namely volumetric CT data, pre- and intra-operative laryngoscopic video and patient-specific CFD simulation that shows the air flow necessary for

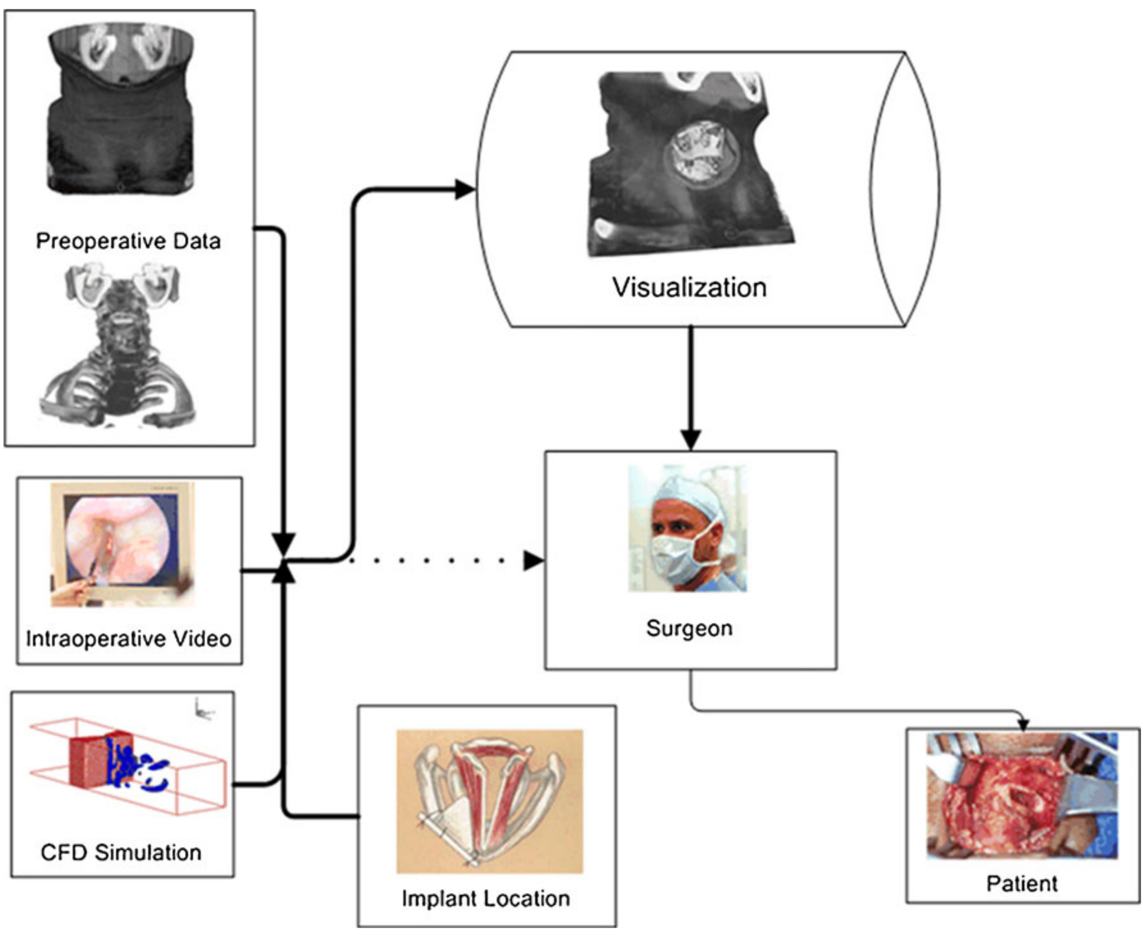


**Fig. 10** Sample workflow diagram of our visualization approach applied to Medialization Laryngoplasty

phonation. The surgeon has to mentally fuse all the information and decide the corresponding position on the patient to successfully place the implant at the optimal position.

Our approach can be integrated into this workflow to improve the efficiency and success rate (Fig. 10). Instead of all the information normally available to the surgeon presented separately (denoted by a dotted arrow in Fig. 10), the surgeon would have the ability to visualize all of the available datasets in an integrated manner. By displaying the anatomical information using magic lens-based data exploration and via volume editing, it is possible to see the relative position of the vocal fold with respect to the laryngeal cartilage, which is normally not visible to the surgeon. Similarly, by fusing laryngoscope and external video camera images with the volume data the surgeon can have better spatial awareness of what he sees in the videos, and can have improved mental registration to help locate target locations on the patient more effectively. The flow simulation data can be incorporated to detect areas with lower flow velocity or irregular flow for phonation. An interactively adjustable implant can be included into the visualization and can serve as a guideline during the surgery for optimal implant placement.

## Conclusion

In this paper, we described a visualization system that integrates multiple modalities for surgical applications. By interactively changing the rendering parameters in a user-selected local volumetric region, effective, and intuitive visualization of different datasets is possible. As examples of our approach, modalities such as computer simulations and real-time camera images, which are conventionally not integrated into medical visualization, were used to improve overall understanding of the relationships between datasets. We believe these visualization techniques can improve the effectiveness of surgical planning and image-guided surgery procedures.

The approaches were developed in close collaboration and with feedback from surgeons and radiologists. Even though we have conducted user studies to evaluate the effectiveness of some of our interaction methods, further investigation of our methods with medically trained users would be very beneficial. We are currently planning on such studies to test the effectiveness of magic lens interfaces and providing the ability of volume editing, as well as if integrating multiple modalities eases the mental fusion process. These kinds of systematic studies about volumetric visualization and corresponding interaction methods are crucial for the increased use of 3D visualization in medical practice.

Our approach is flexible in terms of addition of different modalities and rendering styles depending on the application domain. As a medical application domain, we believe blood flow is a good candidate for our approach to integrate either acquired (Doppler/contrast-enhanced ultrasound, contrast CT/MRI) or calculated (cardiac strain, blood flow simulations) flow datasets with traditional anatomical volume information.

**Acknowledgment** This work was supported by NIH grant R01 DC007125 to develop computer-based tools for medialization laryngoplasty. The authors would like to thank Dr. Steven Bielamowicz for his help and informative discussions for application of our methods to otolaryngology, and Dr. Rajat Mittal and Qian Xue for providing the CFD simulation datasets.

## References

1. Correa CD: Visualizing what lies inside. SIGGRAPH Comput Graph 43(2):1–6, 2009
2. Luo H, Mittal R, Zheng X, Bielamowicz S, Walsh R, Hahn J: An immersed-boundary method for flow–structure interaction in biological systems with application to phonation. J Comput Phys 227 (22):9303–9332, 2008
3. Perrin DP, Vasilyev NV, Novotny P, Stoll J, Howe RD, Dupont PE, et al: Image guided surgical interventions. Curr Probl Surg 46 (9):730–766, 2009
4. Peters TM, Cleary K: Image-guided interventions: technology and applications. Springer, New York, 2008
5. Peters TM: Image-guidance for surgical procedures. Phys Med Biol 14:R505, 2006
6. Kruger J, Schneider J, Westermann R: ClearView: an interactive context preserving hotspot visualization technique. IEEE Trans Visual Comput Graph 12(5):941–948, 2006
7. Svakhine N, Ebert DS, Stredney D: Illustration motifs for effective medical volume illustration. IEEE Comput Graph Appl 25(3):31–39, 2005
8. Bier EA, Stone MC, Pier K, Buxton W and DeRose TD: Toolglass and magic lenses: The see-through interface. In: Proceedings of the 20th annual conference on Computer graphics and interactive techniques, Anaheim, CA, 1993, pp 73–80
9. Wang L, Zhao Y, Mueller K and Kaufman A: The magic volume lens: an interactive focus+context technique for volume rendering. pp 367–374
10. Viega J, Conway MJ, Williams G and Pausch R: 3D magic lenses. In: Proceedings of the 9th annual ACM symposium on User interface software and technology, Seattle, Washington, United States, 1996
11. Bruckner S and Gröller ME: VolumeShop: an interactive system for direct volume illustration. In: IEEE Visualization, 2005, pp 671–678
12. Burger K, Kruger J, Westermann R: Direct Volume Editing. IEEE Trans Visual Comput Graph 14(6):1388–1395, 2008
13. Kirmizibayrak C: Interactive volume visualization and editing methods for surgical applications. Ph.D. Dissertation, Department of Computer Science, The George Washington University, Washington, DC, 2011
14. Zhang Z: Iterative point matching for registration of free-form curves and surfaces. Int J Comput Vision 13(2):119–152, 1994
15. Gerasimov P, Fernando R and Green S: Shader Model 3.0: using vertex textures, white paper DA-01373-001_v00, NVIDIA, 2004
16. Scheuermann T and Hensley: Efficient histogram generation using scattering on GPUs. In: Proceedings of the 2007 symposium on Interactive 3D graphics and games, Seattle, Washington, 2007, pp 33–37
17. Li W, Fan Z, Wei X and Kaufman A: GPU-Based flow simulation with complex boundaries. GPU Gems II, 2003
18. Yim Y, Wakid M, Kirmizibayrak C, Bielamowicz S and Hahn JK: Registration of 3D CT data to 2D endoscopic image using a gradient mutual information based viewpoint matching for

image-guided Medialization Laryngoplasty, JCSE 4(4):368–387, 2010
19. Studholme C, Hill DLG, Hawkes DJ: An overlap invariant entropy measure of 3D medical image alignment. Pattern Recogn 32(1):71–86, 1999
20. Johnson R, O'Hara K, Sellen A, Cousins C and Criminisi A: Exploring the potential for touchless interaction in image-guided interventional radiology. In: Proceedings of the 2011 annual conference on Human factors in computing systems, Vancouver, BC, Canada, 2011, pp 3323–3332
21. Kirmizibayrak C, Radeva N, Wakid M, Philbeck J, Sibert J and Hahn J: Evaluation of gesture based interfaces for medical volume visualization tasks. In: Proceedings of the 10th International Conference on Virtual Reality Continuum and Its Applications in Industry, Hong Kong, China, 2011, pp 69–74
22. Helferty JP, Sherbondy AJ, Kiraly AP, Higgins WE: Computer-based system for the virtual-endoscopic guidance of bronchoscopy. Comp Vis Image Understand 108(1–2):171–187, 2007
23. OpenNI Organization. "OpenNI User Guide," 05/01, 2011; http://www.openni.org/documentation
24. Jin G, Baek N, Hahn JK, Bielamowicz S, Mittal R, Walsh R: Image guided medialization laryngoplasty. Computer Animation and Virtual Worlds 20(1):67–77, 2009
25. Bielamowicz S: Perspectives on medialization laryngoplasty. Otolaryngolog Clin North America 37(1):139, 2004
26. Luo H, Mittal R and Bielamowicz SA: Analysis of flow-structure interaction in the larynx during phonation using an immersed-boundary method. J Acoust Soc Am vol. 126:no. 2, 2009