This article has been accepted for publication in IEEE Transactions on Visualization and Computer Graphics. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2026.3686821

Designing a Collaborative Immersive Visualization System for Radiation Treatment Planning Teams

Kiet Tran*, Matthias Broske*, Michael G. Herman, Victoria Interrante, Evan Suma Rosenberg, Daniel Mundy**, and Daniel F. Keefe**

***Abstract*—We present a visualization design study of creating a collaborative virtual reality (VR) system for radiation treatment planning, with an emphasis on proton therapy. The goal is to support teams of dosimetrists, physicians, and medical physicists as they review and compare multiple possible patient-specific treatment plans, which requires analyzing complex 3D spatial relationships between a radiation dosage volume and anatomical structures. The approach is a novel combination and refinement of interactive visualization techniques including: networked multi-user immersive visualization, interactive volume rendering and slicing with 3D widgets and gestures, superimposed surface rendering with GPU-accelerated curvature-directed lines, smart cursors, teleporting, and avatars. These features are integrated within a workflow that supports three complementary modes of visual data comparison (juxtaposition, interchangeable, and explicit encoding). Results and feedback from multi-year iterative development with users and a summative field deployment in the form of a mock plan-review meeting reveal several advantages relative to current clinical practice and suggest directions for future work.**

***Index Terms*—Collaborative virtual reality, comparative visualization, volume rendering, radiation treatment planning.**

I. Introduction

NEARLY two million new cancer cases occur in the US each year [56], and approximately 500,000 are treated using radiotherapy [9]. Each individual radiation therapy treatment plan can take up to 1-2 weeks to define [13] because of the complex challenges in selecting the optimal treatment parameters to ensure successful delivery of prescribed irradiation to the target regions while minimizing the irradiation of uninvolved tissues. Immersive virtual reality technology has tremendous potential to increase both the efficiency and effectiveness of the treatment planning process by facilitating the understanding of, and communication about, the complex 3D dose distributions predicted by multiple plans across myriad regions within a patient's body.

This paper describes our interdisciplinary team's collaborative process and reflections from a three-year project to explore the potential of immersive visualization to address challenges faced by medical researchers and clinicians at the Mayo Clinic's Proton Radiation Treatment Center in Rochester, Minnesota, USA. For the visualization and virtual reality (VR) research communities, this paper can be viewed as a **design study**, following the problem-driven research methodology described by Sedlmair et al. [55]. Consistent with this framing, this paper describes: 1. analyzing a specific real-world problem faced by "domain experts" (a term used in the visualization literature to describe expert collaborators who face complex data understanding problems); 2. characterizing this problem through an abstraction into tasks and data to establish the requirements by which a design proposal should be judged, 3. designing an interactive software tool to address this problem, including discussion of alternative designs, 4. validating the design through iterative development with domain experts and a field deployment to provide evidence that the tool is useful to the experts, 5. a reflection on lessons learned that can improve current guidelines for effective immersive visualization.

* The first two authors share the first-authorship position.
** The last two authors share the senior-authorship position.
Kiet Tran, Matthias Broske, Victoria Interrante, Evan Suma Rosenberg, and Daniel F. Keefe are with the University of Minnesota. E-mail: tran0563@umn.edu, tias@broske.com, {interran|suma|dfk}@umn.edu.
Michael G. Herman and Daniel Mundy are with the Mayo Clinic. E-mail: {Herman.Michael|Mundy.Daniel}@mayo.edu.

II. Background on Radiation Treatment Planning

Radiation treatment planning is difficult because it is impossible to achieve an ideal plan, delivering a full dose of radiation to cancerous tissues and zero dose to healthy tissues. Proton therapy is more precise than conventional therapy based on X-ray beams and is, therefore, particularly suitable for treating tumors near critical organs or in pediatric patients, where minimizing radiation exposure to healthy tissues is crucial. However, myriad uncertainties, from minute variations in patient positioning to shortcomings in the precision with which the biological effects of the radiation can be estimated via different calculation methods, contribute to the complexity of the process. The treatment planning process cannot be completely automated. An expert team of radiation oncologists, dosimetrists, and medical physicists must consider numerous complicated factors to make an optimal choice among possible plans whose efficacy scores may be numerically equivalent.

Radiation oncologists—physicians who specialize in radiation therapy—are mainly responsible for the prescription, approval, and supervision of the treatment. Dosimetrists are responsible for creating an optimal treatment plan that delivers the prescribed dose to the tumor while minimizing dose to surrounding healthy tissues using advanced treatment planning software tools according to the radiation oncologist's prescription. Medical physicists develop planning tools and strategies and are responsible for quality assurance of the treatment plan, including dose calculation accuracy and safety review [54].

One essential aspect of the team's process is the plan review meeting. During these sessions, the team discusses various treatment plans, evaluates and compares their effectiveness and quality, and makes any necessary adjustments.

Authorized licensed use limited to: University Of Minnesota Duluth. Downloaded on June 23,2026 at 16:38:47 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Visualization and Computer Graphics. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2026.3686821

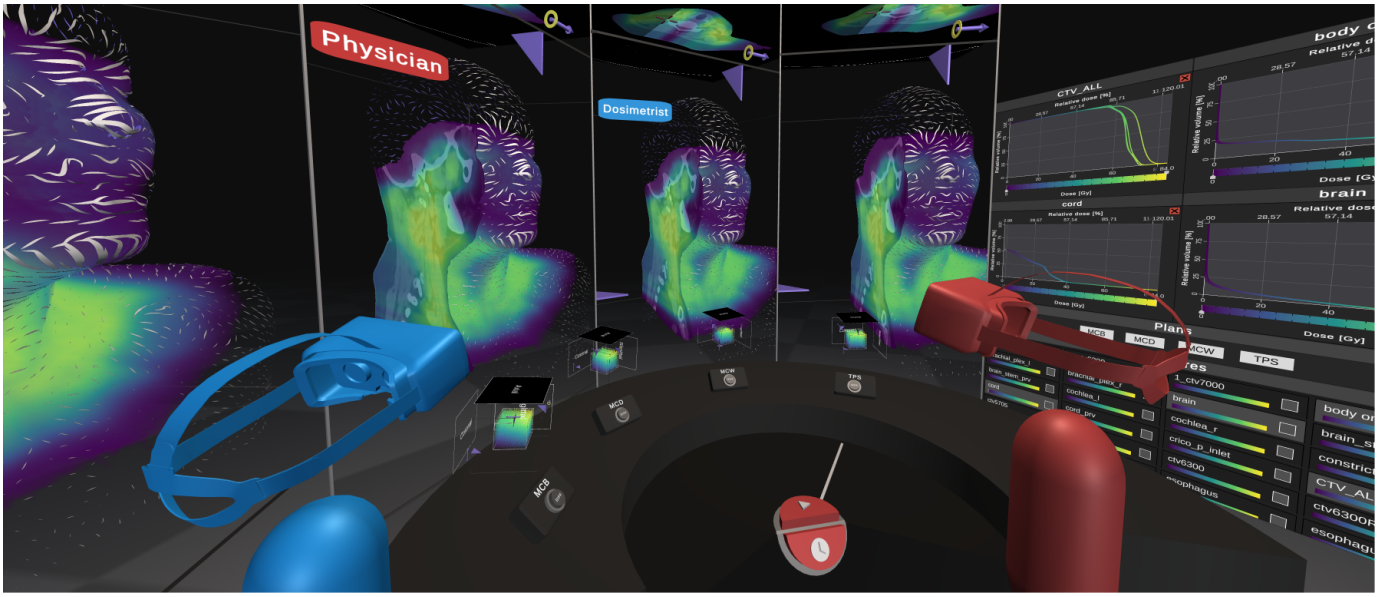

Fig. 1: A medical physicist is joined by two collaborators during a radiation treatment plan review meeting. Four volume rendered treatment plans are visualized for comparison inside portals arranged in a semicircle. The virtual environment includes a bimanual user interface and a variety of widgets for interactive data exploration, including linked 2D *Dose-Volume Histogram (DVH) plots* (far right). The semicircular control table in the center holds four *Plan Display Widgets* for toggling plans and visualization features and a *Plan Clock* for controlling several complementary visual comparison modes. A *Structure Selection Menu* is visible in the bottom right, and a *Bookmarks Menu* is located just off screen to the left.

III. Problem Characterization and Abstraction

Consistent with the design study methodology [55], we present a characterization of the key domain expert problems synthesized through the collaboration and map them to techniques and challenges studied in visualization and VR. For the current data pipeline, we learned that good tools and processes are already in place to prepare medical imaging (CT and/or MRI) volume datasets for each patient, including segmenting the 3D boundaries of treatment regions and anatomical structures to avoid. Then, commercial software is used to develop a first treatment plan based on the prescription from a radiation oncologist. The *prescription* specifies how much radiation to deliver to segmented 3D structures, and the *plan* specifies how the mechanical device that delivers the radiation should be oriented and activated to deliver the radiation. In proton therapy, the increased precision that is possible adds degrees of freedom to the search space for an optimal plan.

The next step is understanding how well the plan meets the prescription and revising and developing alternative plans to identify the best choice, and this is where we identified the most potential benefit for advanced visualization. This is an iterative and collaborative team process that can last one to two weeks and requires close communication and coordination between team members serving in the three key roles described earlier: radiation oncologist, dosimetrist, and medical physicist. Learning through many discussions and demonstrations, we were able to abstract the team's needs during this stage into three specific requirements for a visualization system.

Requirement 1 (R1). The first is accurately visualizing spatial relationships between the dose and patient's anatomy, including structures identified as targets and organs-at-risk. Visualization researchers will recognize this human visual perception problem: how do we optimize 3D computer graphics renderings to help humans accurately judge how one 3D form fits inside or overlaps another. Early research suggests perspective-tracked, stereoscopic rendering helps and provides guidelines on rendering styles that enhance perception of 3D shape [25], [26], [39], but open research questions include: 1. how to translate insights from these early perceptual studies into a real-world context where many structures of interest are likely to be displayed simultaneously, and 2. how to leverage modern, GPU-accelerated VR rendering.

Requirement 2 (R2). The second requirement is supporting collaboration. This requires overcoming some technical challenges around synchronizing data and displays. However, the greater challenge is to design user interfaces and visualizations to support natural collaborative workflows. Prior research includes VR techniques that make it possible to manipulate individual views while maintaining awareness of the group [33] or to simulate face-to-face interactions for training medical teams [30]), but our focus is slightly different. We expect team members' roles to be asymmetric, and the most important task to support is group discussions that regularly reference spatial relationships. Rather than simulating the real-world for clinician training or patient education, our solution might benefit from features that are not possible in real-life, such as each participant viewing the data from a different perspective.

Requirement 3 (R3). The third requirement is supporting comparative data analysis. This is a critical topic in visualization research that has been well-studied for 2D information visualization [19] but less so for 3D and 4D spatial data [31], [42]. Kim et al. [31] identify four fundamental approaches to support visual comparisons of 3D data: Superimposition, Juxtaposition, Interchangeable (animation over time), and Explicit Encoding. They suggest "hybrid" techniques may be particularly useful, but this is an understudied area.

Authorized licensed use limited to: University Of Minnesota Duluth. Downloaded on June 23,2026 at 16:38:47 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Visualization and Computer Graphics. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2026.3686821

IV. Related Work

A. Related to R1: VR and Perception in Radiation Treatment

Preim et al. provide an introduction to visual computing for radiation treatment planning in their recent book [50]. Specific to VR and AR technologies, Su el al. applied VR to training medical practitioners [59] and Bannister et al., provide a more recent example [3]). VR and AR have also been used to educate patients about radiotherapy and/or to facilitate relaxation (e.g. recent examples [22], [62]).

Our area of interest is using VR for spatial analysis during dose plan review. A 2019 state-of-the-art report [54] identified eight prior works using VR, AR, or a holographic display in this way [5], [14], [16], [17], [48], [59], [61], [63]. One of the eight prior systems was evaluated in a multi-site study, and the authors found, in part, that the radiotherapy plans created using 3D visualization were "preferred in a majority of cases over those developed with 2D displays" [14]. Our work builds on these with several differences. First, seven of the eight prior works are brief reports of three to five pages each that do not describe or evaluate the computer graphics and user interface design decisions in detail. The eighth describes the multi-site study in detail but does not make a visualization or user interface contribution. We present a detailed technical description of the system and all major design decisions backed by technical references. We even explain the design rationale for the VR cursor used in our implementation. Second, our work is a modern application with features like side-by-side interactive volume rendering in VR that are only made possible with modern hardware.

From a visual data analysis standpoint, the fundamental task required for radiation treatment planning is understanding how one volume or isosurface (representing the radiation dose) fits within or overlaps with a second volume or segmented surface (the patient's anatomy) [25]. To date, the most perceptually optimal rendering techniques for facilitating such surface-inside-a-surface spatial judgments rely on non-photorealistic techniques that render at least one surface with a special texture that is completely opaque in some regions and completely transparent (see-through) in others [2], [25], [26]. The best textures are computed so that the opaque regions are lines that follow the gradient of surface curvature. This has two perceptual benefits. First, the curvature-directed lines make the shape of these complex 3D surfaces easier to understand. Second, the fact that much of the surface is transparent makes it easier to see the second surface inside the first [26]. While this technique has previously been applied with textures generated offline, we introduce a GPU-based implementation.

B. Related to R2: Collaboration and Interaction in VR

Synchronous collaborative systems have been designed for visualization with desktop displays to tabletop systems (e.g. [11], [28], [58]), and the research has led to novel strategies to support asymmetric roles within the group (e.g. [38]), converging and diverging workflows (e.g. [60]), and more (e.g. [10]). In immersive visualization environments, collaborators often use individual displays; thus, the state of the visualization must be synchronized across the displays. Another challenge (and opportunity) is that collaborators can take on the same perspective of the scene, as if everyone is sitting in the same chair, or take completely different perspectives that may be meters or kilometers away from each other in the virtual world [64]. Representing the current state of the team is, therefore, an important design consideration, often addressed using avatars [4], [57]. Collaborative VR systems have been developed to support users with symmetric roles (e.g. [33], [34]). Another paradigm is for one collaborator to be an instructor/guide while others follow (e.g. [12], [30], [36]). Our system supports asymmetric collaborative workflows through several features, e.g., we support both gesture- and widget-based manipulation of the visualization volumes. Gesture-based interactions tend to be faster but are not self-revealing, so they are often preferred by frequent tool users whereas widget-based techniques are preferred by novice and less-frequent users [37].

The smart-cursors, widgets for system control and data exploration, and linked 2D and 3D visualizations in our implementation are well informed by research on spatial user interfaces and immersive visualization and analytics [15], [21], [35], [43], [49]. We consider the 3D user interface representative of the state of the art.

C. Related to R3: Visual Comparisons of 3D Data (in VR)

Comparison is fundamental to data analysis and is an essential task to support with visualizations across many fields, including medicine [18], [19]. While comparison is often explicitly supported with 2D visualization systems [42], visual techniques designed specifically to support comparative analysis are less common when visualizing 3D spatial data, especially within VR environments.

According to a survey on comparison techniques for spatial 3D and 4D data [31], only 2 of the 41 systems identified utilize VR. Extending the classification introduced for 2D visual comparisons [19], the survey classified visual comparison techniques into four fundamental approaches: Superimposition, Juxtaposition, Interchangeable (animated over time), and Explicit Encoding. The authors call specifically for new research on combining multiple approaches within the same application. Our approach is a hybrid, combining Interchangeable, Juxtaposition, and Explicit Encoding and drawing inspiration from multiple prior VR systems. First, like the Bema system, which uses an Interchangeable approach with an interactive timeline control to animate between the different phases of an archaeological site [32], our system includes interactive (and automatic) animations between different radiation plans. The theory is that the animated movement only in regions of difference draws the eye to these areas [41]. Second, our system uses a Juxtaposition approach inspired by the Worlds-in-Wedges [46] technique, where multiple data instances are arranged around the user in pie-slice shaped volumetric portals that support linked or individual 3D interactions. We demonstrate how Explicit Encoding can be integrated with Worlds-in-Wedges by interactively selecting one pie-slice to act as the basis for comparison with the others.

Authorized licensed use limited to: University Of Minnesota Duluth. Downloaded on June 23,2026 at 16:38:47 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Visualization and Computer Graphics. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2026.3686821

V. Detailed Technical System Design

Fig. 1 provides a high-level orientation to the major components of the system, and this section describes the visualization algorithms and user interface in detail. The system runs on current generation VR, AR, and MR head-worn displays and assumes users provide input via two handheld controllers, each with one trigger button and one grip button. In addition, for the dominant hand, the interface also utilizes the small joystick control and "A" and "B" buttons found on most current commodity VR controllers. The application is implemented in C# using the Unity game engine with the Built-In Rendering Pipeline and a variety of custom user interface techniques and compute and rendering shaders as described in the next sections. A desktop-level graphics card is required; so, a tethered desktop-streaming mode is used for headsets with integrated processors, such as the Meta Quest devices. When users enter the virtual environment, they select the dataset to load using a menu, and this creates a preliminary visualization similar to that shown in Fig. 1.

A. Interactive Widgets and 3D Data Exploration

Many of the interactions utilize a Smart Cursor and a pointing metaphor. Our implementation is based upon the Bendcast technique [51] with the added feature that the laser pointer can bend to snap to the closest selectable object. The technique has been described [37] as the pointing analog of the 3D Bubble Cursor [20]. The cursor is controlled via the 6 degree-of-freedom tracking input from the dominant hand. Visual feedback changes dynamically as the cursor snaps to different widgets as seen in the accompanying video.

View/volume manipulation. Volume rendering cubes can be adjusted via two complementary techniques. The first uses 3D widgets displayed on the volume cube for scaling, rotating, and translating the cube. These are activated by snapping the smart cursor to the appropriate widget then holding the trigger button. To increase precision and reduce fatigue, all three widgets are controlled through small, relative movements of the controller with visual feedback provided in both the volume visualization space and the small control space. This control space is facilitated by a center of movement close to the controller, calculated as a set distance based on the direction vector from the widget to the volume rendering cube's center. The second technique, designed for users with more VR expertise, is a bimanual manipulation similar to SmartScene [23] and Spindle [44]. Here, translation is accomplished by holding a single grip button (either on the left or right hand). Holding the second grip button then activates a combined rotation and scaling mode. In all cases, when multiple volume cubes are displayed, their transforms are linked so that manipulating one causes the others to update.

Slicing Planes. The volume cube includes CT slicing planes aligned with the sagittal, axial, and coronal anatomical planes, and these can also be adjusted interactively by grabbing them with the smart cursor and then moving within a small local control space. While grabbing a slicing plane, tapping the thumbstick in the left or right direction changes the half of the volume rendering that is clipped from view. A crosshair is drawn at the intersection of the three planes, and this is also an interactive widget that can be grabbed with the smart cursor. Moving the crosshair in one plane causes the CT slices for the other two planes to move accordingly.

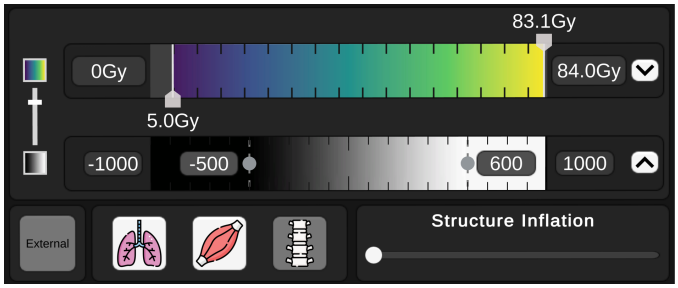

Fig. 2: The Visualization Shortcuts Palette floating above the non-dominant hand includes the Dose Range Slider (colored gradient), CT Range Slider (gray-scale gradient), and toggles for rendering of anatomical context (bottom icons).

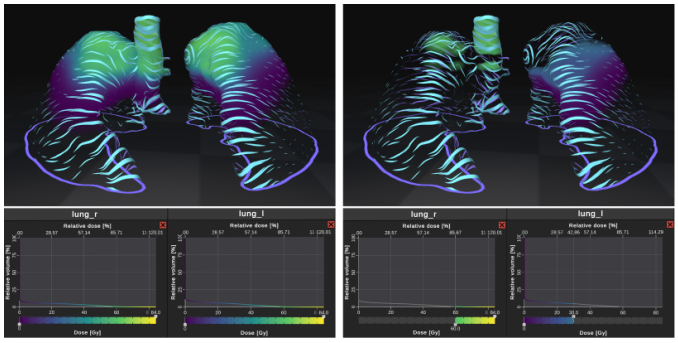

Fig. 3: Left: Each selected structure is linked with a DVH plot, and each DVH plot has a dosage range slider. Right: Editing the dosage range alters the transfer function for the linked structure.

Visualization Shortcuts Palette. The most important visualization and system controls are available for quick access as virtual palettes that can float above the user's non-dominant hand. A close-up of the palette is pictured in Fig. 2. All of the buttons and sliders are activated using the smart cursor, as demonstrated in the accompanying video.

2D plots and menus. Smart-cursor-based interactions are used to work with all of the menus and 2D plots. These include the Structure Menu (far right in Fig. 1), which is used to toggle visibility of anatomical and radiation structures. An interactive DVH plot is also displayed for each active structure (Fig. 3). The X-axis represents the relative dose in percent or absolute radiation dose in Gray (Gy). The Y-axis represents the absolute or relative volume of the structure that receives at least the dose indicated on the X-axis. Each plot contains one curve for each treatment plan, and clicking the curve toggles the visibility of the corresponding plan in the 3D view. A structure-specific dose range slider is also included, as shown in Fig. 3. Menus are also used to save and load the current state of the visualization.

B. Comparison Modes

In the Juxtaposition comparison mode (shown in Fig. 1), multiple treatment plans are visualized side-by-side. This

Authorized licensed use limited to: University Of Minnesota Duluth. Downloaded on June 23,2026 at 16:38:47 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Visualization and Computer Graphics. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2026.3686821

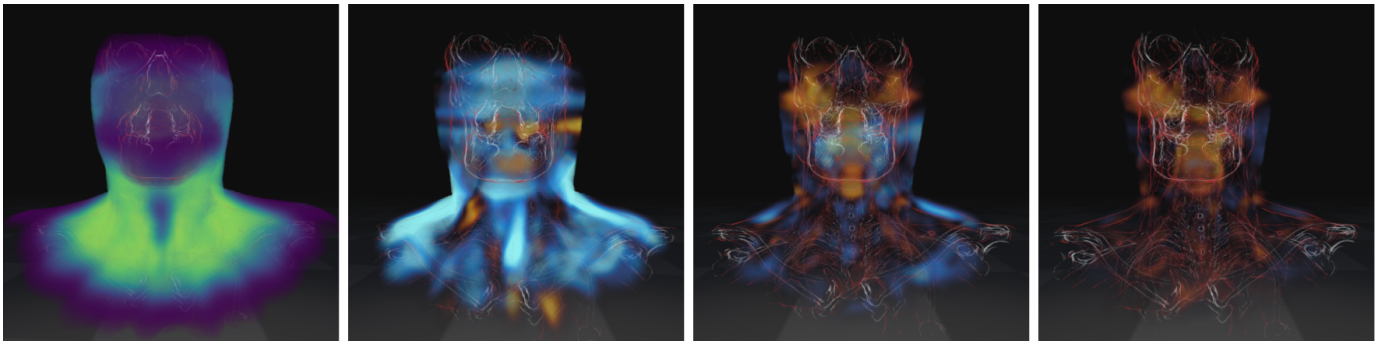

Fig. 4: In Differences (Explicit Encoding) mode, the reference volume (far left) keeps its original purple-green colormap, while the difference volumes employ a divergent blue-yellow colormap, with blue showing cooler dose and yellow showing higher dose.

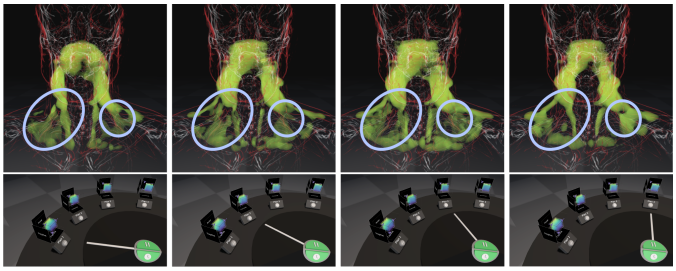

Fig. 5: In Interchangeable mode, all of the side-by-side volume rendering instances collapse into one single instance. This instance animates through the different plans (top images). The Plan Clock on the table indicates which plan is being visualized (bottom images). Here we have circled some areas of subtle differences in the volume that could catch the eye when viewed in VR.

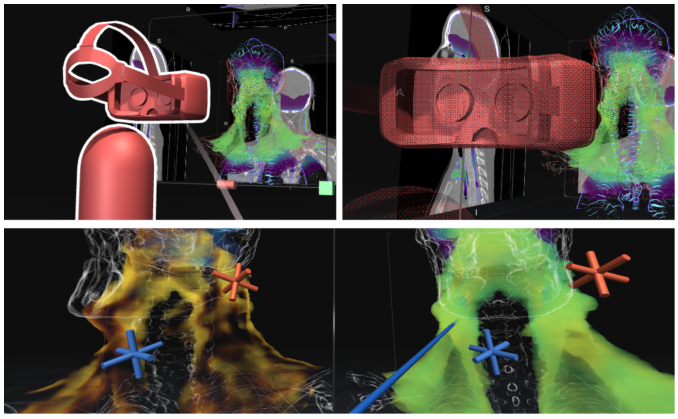

Fig. 6: Top left: Selecting another user's avatar teleports to their view with a gentle animated transition. Top right: Whenever an avatar would block the view of the data, a dithering effect gradually fades the avatar from view. Bottom: User 1 (blue) and User 2 (red) have each placed a crosshair within the volume visualized in the left portal, and the same spatial locations are identified in all other portals.

mode also supports an Explicit Encoding of difference, which is activated by clicking the "DIFF" toggle below the nameplate for one of the plans. This makes the plan the reference for comparison, and the visualization for every other plan is updated to display a voxel-by-voxel difference of the plan minus the reference plan. Fig. 4 demonstrates the effect.

The Interchangeable comparison mode (Fig. 5) is controlled with the Plan Clock widget, which includes an On/Off button on the side of the semicircle table directly facing the user. In this mode, the side-by-side volume renderings are collapsed into one that animates through the individual plans. The Play/Pause button toggles automatic animation. Alternatively, pressing the thumbstick in the direction of the plan you wish to see controls the animation manually. The clock hand always points to the plan that is currently displayed.

C. Collaborative Data Analysis

Each user runs a separate instance of the VR application, and these client applications each connect using the TCP network protocol to a separate server written in C++ for synchronization (whose implementation can be found on the *ivlab* Github page, under the *MinVR3-UnityPackage* repository). All users have individual control of their view via head-tracking. Users can also control their personal view of the data by reorienting the volume data using the widgets or gestural interfaces. Control of the finer aspects of the visualization state are assigned to one "driver" at a time. Control can be passed to a new driver using the System Palette, available on the non-dominant hand, similar to the Visualization Shortcuts Palette.

Each user is represented by an avatar, and it is possible to teleport to another user's view with an animated transition (Fig. 6 top row). Two simple annotation tools (crosshairs and an extendable pointer) facilitate collaborative discussion of spatial features identified in the data (Fig. 6 bottom). The extendable pointer is activated by clicking an icon floating by the non-dominant hand. Its length can be adjusted using the thumbstick and it does not snap to nearby widgets.

D. Interactive Volume Rendering

We employ a raycast-based approach to volume rendering; the final output color for each pixel is determined by stepping along a ray cast into the volume and accumulating color via a transfer function. However, we add several optimizations and application-specific features.

Signed-distance fields. The volume renderer makes extensive use of 3D Signed-Distance Fields (SDFs) to achieve performant volume rendering of multiple proton dosage volumes side-by-side in VR. The SDFs are calculated in a pre-processing step and automatically saved to disk so the calculation is only required once for each new dataset loaded. The first SDF of interest is the Plan Distance Field, which describes the approximate distance from every empty voxel to the nearest non-empty voxel. This SDF is generated for

Authorized licensed use limited to: University Of Minnesota Duluth. Downloaded on June 23,2026 at 16:38:47 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Visualization and Computer Graphics. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2026.3686821

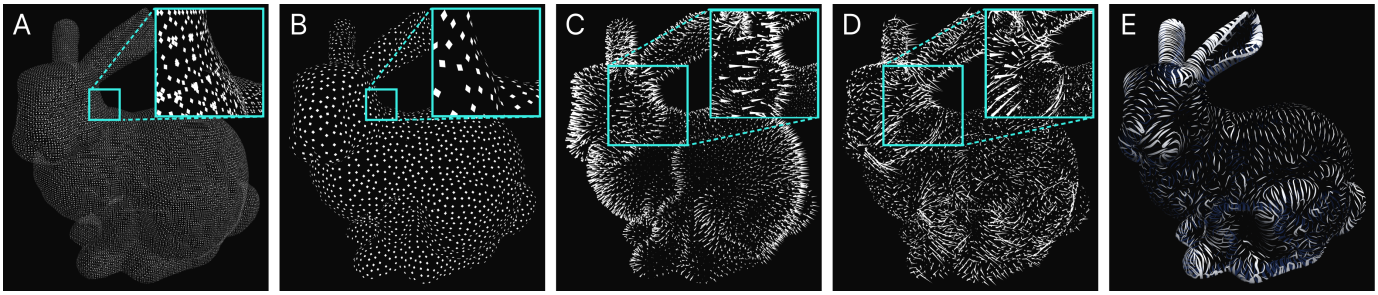

Fig. 7: The GPU pipeline for generating curvature-directed streamlines, applied to an SDF of the Stanford Bunny. Arbitrary points are densely sampled across the entire surface of the SDF (A) and then transformed into a uniform distribution via Poisson disk sampling (B). The gradient is computed at every voxel in the volume (C), which in turn is used to derive the first principal directions and curvatures for each voxel (D). Finally, streamlines are generated on the GPU by tracing along the first principal directions, scaling width and length by first principal curvatures (E).

each plan using the Jump Flood Algorithm. By sampling the Plan Distance Field at each step taken along each ray cast, the volume renderer can not only determine when the current voxel is outside of the dosage and skip all transfer function calculations, but also increase the size of the next step. We found both optimizations to be critical.

Structure aware rendering. The volume rendering is designed to leverage the structure contouring that the treatment team already carefully conducts, and these are also used to pre-process volume data. For each structure of interest, the system expects a sequence of contour loops defining the edges of the structure. During pre-processing, these contours are read into a custom GPU data structure and a compute shader is dispatched to generate a Structure SDF for each structure by determining the distance from each voxel to the nearest contoured loop of points, flipping the sign of the distance if it is inside a loop. These SDFs are also automatically saved to disk.

With no structures selected (default) the volume rendering displays data from all non-empty voxels as determined using the Plan Distance Field. (With a toggle on the Visualization Shortcuts Palette, this default can also be switched to display only the voxels within the special structure named "Body".) When users select one or more structures from the Structures Menu, the rendering changes to only display voxels within these structures. The SDFs of the selected structures are combined with a real-time calculation on the GPU to create a single Combined Structure SDF, which is then applied to the volume rendering as a smooth mask by ignoring the dosage wherever the Combined Structure SDF has a value greater than a specified threshold. In general, a threshold of zero is used; however, positive values have the visual effect of "inflating" the structures, and we have mapped this to a useful interactive control demonstrated in the accompanying video. Squeezing the elastic-feedback trigger button on the dominant hand gradually inflates the structure to see the dosage nearby.

While the Visualization Shortcuts Palette includes sliders that control the global transfer functions, it is also possible to set a custom Structure Dosage Range. This requires assigning each voxel to a structure, and this is accomplished using the Jump Flood Algorithm [52] to create a mask volume where each voxel stores a code identifying the nearest structure. This approach works even for "nearby" voxels when structures have been inflated. In the case where voxels belong to multiple structures, the most recently selected structure is assumed.

Anatomical context. Rendering anatomical context from CT data within the dosage volume also requires some pre-processing. First, a copy of the CT data are downsampled to be the same size as the dosage volume. Second, we compute the gradient at each voxel in the downsampled volume using a simple central difference. The downsampling improves performance by reducing the cache miss rate on the GPU, while pre-calculating the CT gradients avoids redundant calculations each frame during volume rendering.

During rendering, if a given point within the volume has a CT gradient above a specified threshold and the gradient is orthogonal to the view direction within a separate specified threshold, then this is considered a silhouette edge. Since this effect is intended only to add subtle context to the visualization, the renderer simply classifies the location within the volume as bone, soft tissue, or airway based on Hounsfield units, and then blends white, red, or blue color with the dose color for the location. Visibility for these three anatomical layers can be toggled on the Visualization Shortcuts Palette.

Supporting volume-to-volume comparisons. Additional logic is required to support the Interchangeable and Explicit Encoding comparison modes. When Explicit Encoding mode is entered, a compute shader is dispatched to calculate the difference volumes for each plan by subtracting the selected plan from them, and an atomic min/max difference is computed in the process. Although similar functionality could be obtained by pre-calculating these differences, computing the min/max difference every time the primary difference plan is changed or a structure/plan is selected or deselected, makes it possible to indicate the min/max difference for the current state of the visualization by displaying it on the Difference Range Editor, which replaces the Dose Range Slider on the Visualization Shortcuts Palette when the difference comparison mode is active. This widget includes functionality to stretch the bounds of the divergent colormap used in difference mode and to mask any dosage below a specified difference value.

In the Interchangeable comparison mode, the volume renderer samples from two plan dosage volumes at every step, mixing between the dosage values. The Interchangeable and Explicit Encoding modes can be active at the same time.

Curvature-directed structure rendering. The volume ren-

Authorized licensed use limited to: University Of Minnesota Duluth. Downloaded on June 23,2026 at 16:38:47 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Visualization and Computer Graphics. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2026.3686821

dering is combined with a perception-enhancing non-photorealistic surface rendering technique to display the bounds of selected structures. The approach is a modern interpretation of prior research on curvature-directed texturing [27] and generates streamlines procedurally on the GPU following the process illustrated in Fig. 7.

Evenly spaced streamlines improve the perceptual effectiveness of the technique so the first step (Fig. 7 A-B) is to generate a random even distribution of points on the surface to act as seed points for the streamlines. We follow the approach suggested in Bowers et al. [6] to generate a Poisson distribution in real time on the GPU; however, instead of generating our initial large set of random points from triangles, we work directly from the Structure SDF using a compute shader that interpolates neighboring voxel positions to generate points that lie along the isosurface of the structure.

Next, we compute the first principal direction along with the first principal curvature for each voxel in the structure SDF, following the approach described by Interrante [24]. We begin by calculating the gradient of the SDF at every voxel (Fig. 7 C), and then use those gradients to create a 2D orthonormal basis in the tangent plane at each voxel from which we can derive the principal direction (Fig. 7 D) by diagonalizing the second fundamental form.

Our GPU-based algorithm then runs a compute pass over the set of streamline starting points, determining the min and max first principal curvature for just these points. These extreme values are useful for adjusting the visual attributes of the streamlines (e.g., color, width) based on the local curvature.

Finally, a compute shader generates streamlines outward in both directions from each starting point using the fourth order Runge-Kutta method to integrate a path along the first principal direction. The principal direction from the previous step is used to maintain heading, and the traced point is projected onto the surface after each step using the gradient of the SDF and the distance to ensure streamlines remain fixed to the surface of the structure. Streamline length and width is determined by remapping the first principal curvature value of the starting point using the previously calculated min and max values. Since the triangles for the streamlines are generated directly on the GPU, they can be quickly rendered.

This technique is highly dependent on the distance at which the gradients are sampled when computing the first principal direction and curvature: too large of a distance and the surface is likely to be considered uniformly curved, too low of a distance and the slightest of changes in the gradient from voxel-to-voxel will create artifacts in the final image. We pre-compute an "ideal" sampling distance for each structure by calculating the coefficient of variation in first principal curvature values for a series of increasing sampling distances and selecting the distance for which first principal curvatures have a coefficient of variation closest to one. Good default values for all other streamline rendering parameters, such as line length, line width, and line spacing, are calculated as scalar multiples of this sampling distance. However, the application also includes an interactive panel where users can fine-tune the rendering in real time by adjusting the streamline length, streamline width, poisson disk sampling radius, and the distance at which curvature is sampled. The front and back faces of the streamlines are rendered in different colors to further enhance shape perception.

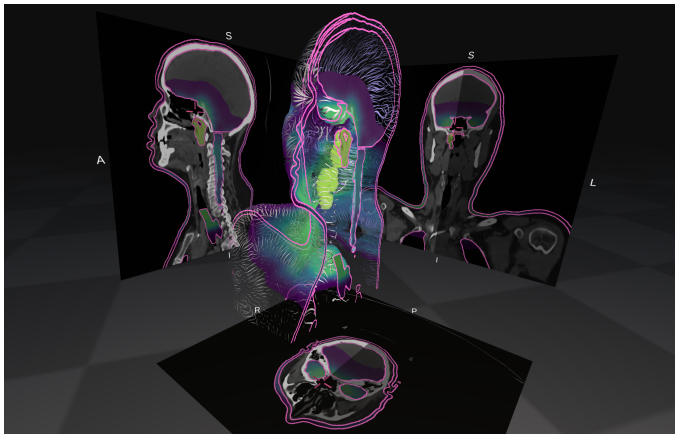

Fig. 8: Three axis-aligned planes are rendered on the far sides of the volume data cube, one each for the sagittal, coronal, and axial planes. The volume is also cut by a sagittal plane. Every plane includes: 1. A uniform thickness outline for all currently selected structures, 2. The dose at that slice rendered in color, 3. CT data at the slice rendered in grayscale.

Slicing plane rendering. Slicing planes can be drawn on the three planes of the volume cubes that are farthest away from the viewer and/or directly within the volume. In both cases, the renderings include uniform thickness outlines around the selected structures and support blending the CT data rendered in grayscale with the dose rendered in color (Fig. 8).

Rendering uniform outlines around the structures is not trivial. A simple threshold based on distance from the structure surfaces in the Combined Structure SDF is insufficient since the Combined Structure SDF is inherently 3D. Instead, the CT slicing plane shader explicitly checks whether the given pixel is within a uniform distance of the structure by searching within its own 2D plane.

For slicing planes, only the pixels that contain structure outlines or visible dosage are rendered to ensure the CT occludes as little of the volume as possible. To correctly handle occlusion, each slicing plane is rendered with an additional shader pass to first record the depth of each visible pixel in the slice. Using these data, the volume renderer is able to ignore pixels occluded by planes in front of the volume and terminate volume traversal early when hitting planes inside of the volume itself. The volume rendering can also be clipped to hide the volume in front of or behind each slicing plane.

VI. Verification via Field Deployment

Now, after describing the technical system design and detailed design choices, we return to the interdisciplinary collaboration and efforts to verify that the system is useful. It is important to note that in a typical design study methodology, there are multiple levels of verification [55]. In our case, the major features, user interface design decisions, and rendering techniques were verified in a lightweight manner during iterative development with multiple cases of anonymized data provided by the two "domain experts" who attended regular project meetings over multiple years to provide feedback.

Authorized licensed use limited to: University Of Minnesota Duluth. Downloaded on June 23,2026 at 16:38:47 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Visualization and Computer Graphics. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2026.3686821

This section describes a complementary, more formal, summative verification via a one-day field deployment. The verification is not a controlled experiment, and it comes with some limitations. However, we believe it provides a useful assessment in the context of this interdisciplinary problem-driven research project. Our goal is to provide documentation of how one radiation treatment team used the system in a challenging scenario that mirrors their normal clinical practice as much as possible given workplace constraints. We acknowledge that because the user feedback comes from one of the co-authors and his colleagues, this creates a demand characteristic [7] that could lead to confirmation bias. Furthermore, the small sample size of expert users may affect the generalizability of the observations drawn from this deployment. To mitigate these limitations, we avoid reporting generic or unsupported claims (e.g., "this software is great", "the user interface is intuitive") and focus on documenting domain-specific discussions and insights. In this reporting, some medical terminology will likely be new to visualization and VR researchers, but we include these details as part of the documentation of how this team of experts used the software to conduct their expert work.

A. Data Preparation

We designed the field deployment scenario to mimic a typical plan review meeting by using anonymized real clinical data from a former patient. To provide an extra challenging case for visual comparisons, the data are from one of the final steps in the planning process, where the team has already converged on optimal angles and settings for radiation delivery and is now comparing four different physics simulations of how the dose will move through the tissues of the body. This means the cross-plan differences visualized should be subtle.

The first (primary) plan is the dose calculated analytically by the commercial treatment planning system (TPS). The other plans are the results of a Monte Carlo (MC) calculation. The first of these, designated MCD, takes into account the biological material in which radiation energy is deposited (muscle, fat, bone, air). The second, MCW, calculates dose assuming that the patient is composed of water of varying density. The third, MCB, approximates the biological effectiveness of the protons used to deliver the dose.

The anatomical data are from a CT image in DICOM format with 292 512x512 pixel 12-bit images and a structure set defining region of interest contours. There are 108 structures including 12 target structures (the tumor and associated margins as defined by the Radiation Oncologist), 68 organ-at-risk (OAR) structures representing the healthy tissues surrounding the targets, and 28 optimization structures created by the dosimetrist to guide the TPS plan optimization algorithm. The plan-specific data come from separate data files, also in the standard DICOM format so they are easily registered to the anatomy. Each file consists of a contiguous 236x136x276 32-bit array of dosage pixel data. Finally, each plan includes summary statistics for each of the 108 segmented structures (total volume; min, max, mean, modal, and median dose received) and tabular data for each of the 108 segmented structures (Dose [cGy], Relative Dose [$\%$], Structure Volume [cm^3]) for creating Dose-Volume Histogram (DVH) plots.

We emphasize that the visualization system can be easily integrated with the team's normal workflow and applied to other treatment plans as needed because all of the data described thus far are generated and saved in standard file formats during the team's normal workflow. To use the new data for the plan review meeting, we simply loaded the files onto three computers (specs detailed under *Apparatus*) and started the visualization software. From the main menu, we selected the folder containing the CT, plan, DVH and structure data files. Since this was the first time the data was loaded on each computer, this initiated the data preprocessing step to calculate the signed distance fields and other derived data needed for visualization. The *preprocessing took 45 seconds to complete* on each computer the first time the new data were loaded into the app.

Computationally, the largest contributor to the preprocessing time (we estimate 90%) is calculating the signed distance field for each structure from its DICOM contour information. Our compute shader implementation on the GPU makes this fast– less than 0.5 seconds per structure, which compares to minutes per structure for a CPU implementation we created during our early development.

B. Methodology

We reserved a conference room at the treatment center for a full day. We setup a three-person collaborative VR environment in the room and also projected a view of the virtual scene using the projector installed in the room. This made it possible for observers to follow and contribute to the discussion of up to three participants wearing headsets. We knew that multiple physicians were interested in participating, and it was likely that at least one would be able to join a mock plan review during the afternoon, but we could not be completely certain because their clinical obligations would be their first priority. To maximize the opportunities for feedback, we, therefore, planned to have all other personnel and systems ready for the afternoon window, and early in the day we ran a tutorial for the medical physicist and dosimetrist.

The tutorial was loosely scripted, starting with demonstrating basic controls for 3D selection and view manipulation and then handing control to the participants who were instructed to try rotating, translating, and scaling the volume visualizations and then moving a slicing plane.

Participants. In the afternoon, we conducted a mock plan review with one physician joining the dosimetrist and medical physicist who participated in the morning tutorial. The medical physicist, who is also a co-investigator on the project and paper co-author, was the local organizer for the day. As a key member of the team, they attended biweekly meetings and helped to shape the design of the tool throughout the three-year project; however, they had never before attempted to use the tool for a plan review. They assumed their typical medical physicist role for the plan review. The dosimetrist and physician had each briefly tried early versions of the tool and provided feedback during two prior site visits. The quotes and discussion reported in this section come from this team of three participants.

Authorized licensed use limited to: University Of Minnesota Duluth. Downloaded on June 23,2026 at 16:38:47 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Visualization and Computer Graphics. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2026.3686821

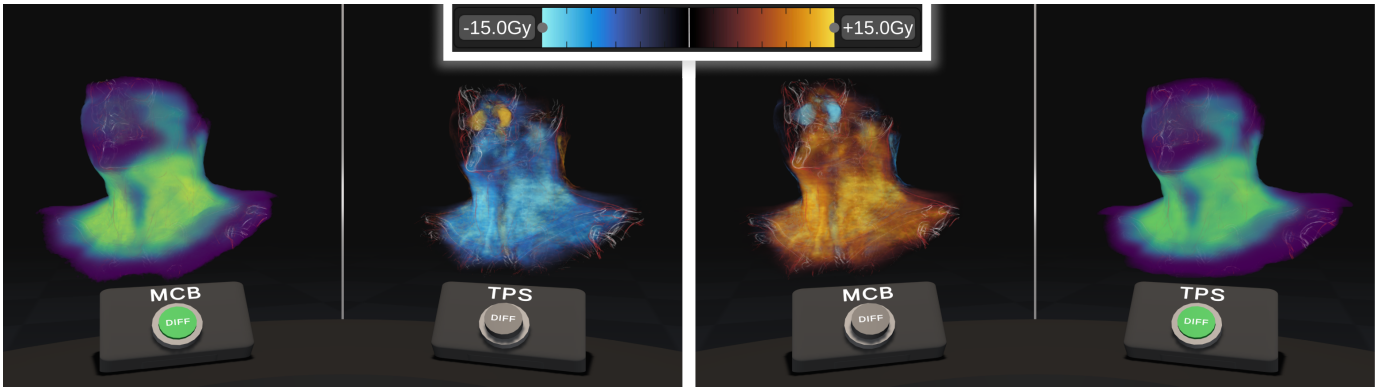

Fig. 9: An observation that confirms an expected trend. Left: When the MCB plan is selected as the reference for the difference comparison mode, the TPS plan updates to display $MCB - TPS$, which evaluates to a negative value for much of the volume. This is expected; TPS should be "cooler" than MCB. Right: Conversely when reference is switched to TPS, it is clear that MCB is "hotter". The cool-warm divergent color map applied to the difference volumes is shown in the top center.

Measures. We encouraged the participants to talk aloud as they worked, recorded the session for note-taking purposes, and took additional written notes as the team worked. When discussing insights from the data analysis, we draw upon conventions used in data visualization research to define *insight* as "an individual observation about the data by the participant, a unit of discovery" [53]. Researchers have demonstrated we can learn about the quality of visualizations by characterizing the insights expressed by users as they work (e.g., via a talk aloud protocol) using criteria, such as the time taken to reach the insight, the domain value, whether the insight leads to a new hypothesis or direction of research, and whether the insight is directed versus unexpected or broad versus deep [47], [53].

Apparatus. Three first-generation Meta Quest VR headsets were used to create the collaborative VR system. Each headset was run in a tethered mode driven by its own dedicated external computer with an Intel Core i7-12700 2.10 GHz processor and NVIDIA GeForce RTX 3060 Ti (8 GB). The three computers were connected to a network switch so they could synchronize the state of the visualization via a local network with one computer running the command-line synchronization server in the background and each VR application acting as a client. In addition to the perspective-tracked, stereoscopic view in each headset, one of the desktops also rendered a non-tracked 3D view of the scene, which was projected on the conference room's screen. All participants were seated.

C. Documentation of the Plan Review Meeting

Training and time to first insight. During the morning training tutorial and after demonstrating how to translate and rotate the volume visualizations, application control was handed off to the dosimetrist who tried this and moving the CT plane widget. Within 1.5 minutes, they began to identify other software features, such as clipping the volume at a slicing plane. As a group, we discussed the interactive dose range sliders; then, in response to a question, we described and tried the toggles for the three views of anatomical context. At eight minutes into the training, the discussion began to turn from understanding features to exploring the data. The conversation included: **D (for dosimetrist):** "This is something you might want to see, where it is enhancing between the two plans." **MP (for medical physicist):** "Yes!" Then the two worked together to reason through how this could be accomplished. They asked a clarifying question about whether by dragging *both* handles on the dose range sliders they could make the blue (low dose regions) disappear entirely and suggested a new feature to show the dose values as a percentage in addition to the absolute dose. Then, they continued to refine the view.

At about 12.5 minutes since starting to work with the tool, there was a clear articulation of a first insight while comparing the TPS dose to the MCB **MP:** "...in this case (Fig. 9 left), the biological plan is the reference and in the physical plan the dose is cooler, which is what you would expect to see." **D:** "Yes, in the righthand side one." **MP:** "Now, if you click the button to set [the physical plan] as the reference... there we go (Fig. 9 right)... Now, we are seeing the biologic dose relative to the [physical] and all of it is orange because all of it is hotter than [the physical], which is what you expect to see." As the team members expressed, the observation that the MCB dose is generally "hotter" than the TPS is correct and expected. This is a broad insight, stating a general trend, and the domain value is low because it is a simple observation of an expected trend. However, confirming the expected is exactly the right place to start. The tutorial continued on for approximately 10 minutes after this, covering the full set of features as they worked with the data visualization.

Beginning the plan review and identifying hot spots. Later in the day, joined by the physician, the mock plan review began with a question, **D:** "Doctor, what would you like to see first in terms of the dose distribution?" The physician replied, **P (for physician):** "Let's walk through the dose levels." The group then spent 4.5 minutes selecting and deselecting structures, adjusting slicing planes, and toggling the visibility of the slicing planes and volume to review dose coverage of the target volumes. There was little discussion at first– it seemed to take some time to get the right view to start the discussion or possibly to determine how to re-imagine the typical plan review process with 3D visualization. Then, the dosimetrist said, "There we go," and the discussion honed in on determining how much of the volume was receiving 70 Gy of radiation as supported by dragging the dose range sliders

Authorized licensed use limited to: University Of Minnesota Duluth. Downloaded on June 23,2026 at 16:38:47 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Visualization and Computer Graphics. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2026.3686821

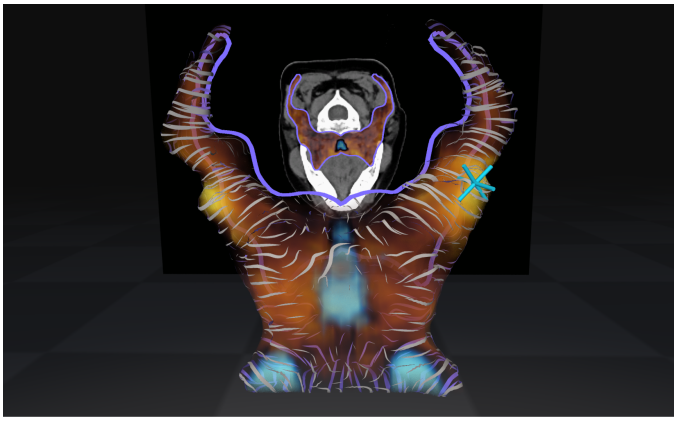

Fig. 10: The team investigated two hot spots that are clearly visible in this 3D structure. During discussion a 3D crosshair marker (cyan color) was placed within the right side of the volume to highlight a point of interest.

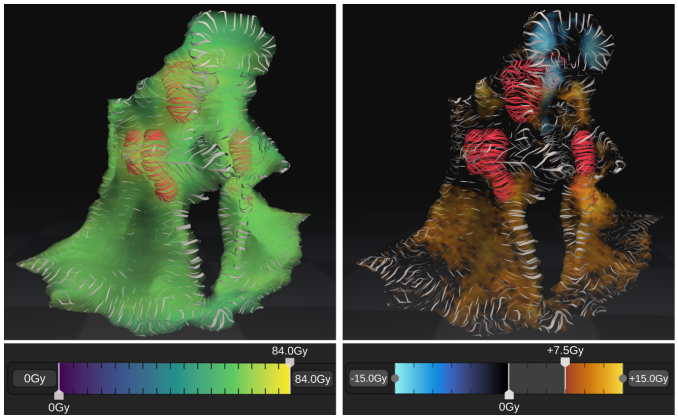

Fig. 11: With the TPS dose as reference (left), we can see the MCB dose (right) includes hotter and cooler regions. Focusing in on the regions that are hotter by 7.5Gy or more, we see yellow and orange areas of bioenhancement, not just around the ctv7000 structure depicted with red lines but also through much of the larger ctv5705 structure depicted with white lines.

to display just the relevant portion of dose.

At six minutes into the meeting, the discussion characterizes how the three communicated about the spatial features they observed, both verbally and by pointing with the virtual cursors. **MP:** "So, now we are looking at 67 to 76... this would be a great view to look at coverage. Because, we would be able to say we have the whole target easily getting at least 95%." **P:** "Yes, and can you localize on the point, like those two posterior points that are 6700?" [pointing with smart cursor] **D:** "You're saying these two little–the islands, right there." [placing their cross-hair cursor while speaking] **P:** "Yes, how do we [focus on / zoom in to] that?" **D:** "Let me grab the axial plane..." **P:** "Yes. Keep going [sliding the plane through the volume]." **MP:** "There it is." **P** "Yes. Why are those there?"

The data exploration had revealed two "hot spots", small regions with higher than expected dose, and the three were beginning to explore in more detail. They toggled the visibility of several structures within the dataset, including one of the target volumes, CTV63. The visualization they created is similar to that in Fig. 10. **MP:** "I appreciate that in the 3D, without having to scroll up and down through the 2D, you can [see the hot spots] right away." **P:** "Yes." **MP:** "Those hot spots are right there at the bottom of that lower-dose CTV."

The discussion turned to suggestions for new features, including complementary linked 2D views that face the viewer separate from those surrounding the 3D volume. This would help to contextualize the new 3D visualizations relative to the 2D images that are used in the current process. This discussion occurred as the dosimetrist continued to explore and move slicing planes through the plan. Then, the physician said, "Keep going. You were going from posterior to anterior... And, then can you turn on the structure in 3D? Let's look at the constrictors or the cricopharyngeal inlet in this view... See, there are areas anteriorly that jump out, right here." [pointing] **MP:** "A little more dose than you expect?" **P:** "Yes." Participants told us that it is not that these hot spots would have been missed in the traditional 2D planning process, but with the 3D visualization they were identified more quickly and easily, specifically they said they were "immediately evident".

A new way to look at bioenhancement. Another period of rich discussion occurred 5 minutes later, or about 20 minutes in total into the meeting. The team was comparing the TPS dose to the MCB with an explicit encoding of dose difference, similar to the visualizations in Figs. 4 and 11. The discussion began with the medical physicist stating an expected trend– they expect to see most of the bioenhancement (visualized as yellow-orange regions in Fig. 11 (right)) on the edges of the high-dose treatment volume (visualized as green in Fig. 11 (left)). The physician then states the unexpected trend that can also be seen in the figure, "but some of it is outside of that, [as shown] in the coronal [slice], right?" **MP:** "Yes." The team spent several minutes using the visualization to explore these regions of high absolute difference in dosage. The current clinical tools do not provide a volumetric visualization *or* calculate an absolute difference across the *full-volume*. So, much of the discussion focused on what the difference visualization reveals and how this new information might be used in the planning process. At one point, the physician summarized, "What is happening in this plan that jumps out that we [would not] see without that [difference comparison mode] is that ... it's biologically hot but we [would not notice this as quickly] because it's biologically hot in the lower dose volume." **MP:** "You are right." The treatment plans include many sub-volumes that often overlap, and what the team found is that the *difference* between the two methods of calculating dosage (TPS and MCB) was relatively high in a volume where the *absolute* dosage was low. These dose variations are subtle and may not be medically significant, but the example demonstrates that the tool supports insights that are unexpected and suggest a deep understanding of the data visualized. This insight required observing trends across multiple structures and plans then synthesizing this new information with expert knowledge of the physics and treatment process.

After the mock plan review meeting, the participants invited several of their colleagues to try the tool. Physicians and dosimetrists who had seen earlier versions of the software (e.g., before the networked collaboration was introduced, before the visual comparison modes) remarked that the tool had advanced significantly. We heard, "this is now beyond a

Authorized licensed use limited to: University Of Minnesota Duluth. Downloaded on June 23,2026 at 16:38:47 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Visualization and Computer Graphics. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2026.3686821

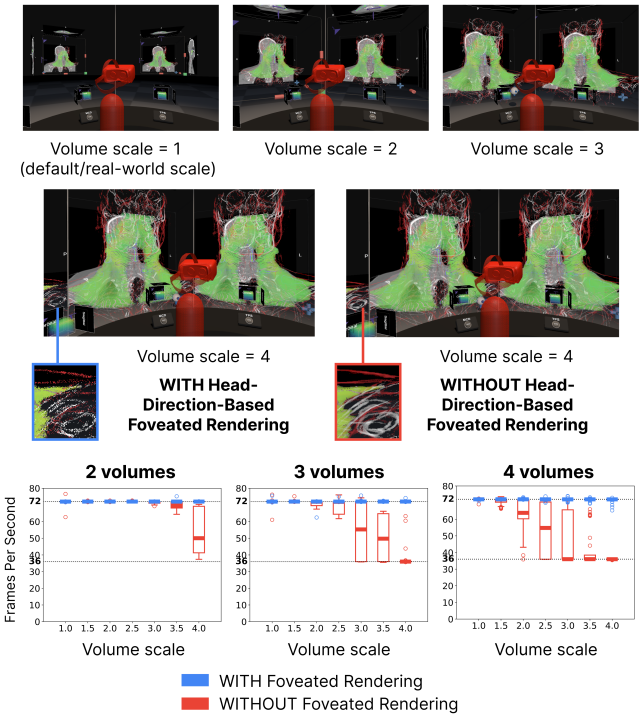

Fig. 12: A comparison of rendering quality and performance with and without the foveated rendering strategy.

proof of concept." A number of future refinements and applications to other clinical scenarios were also suggested. Most of the refinements are for small changes, such as including percentage in addition to absolute values on the dose slider.

D. Evaluating Performance and Foveated Rendering

During the mock plan review scenario, with the three Meta Quest 1 devices running in tethered mode and the maximum 1440x1600 resolution per eye, the hardware capped the framerate at 72 Hz, and the rendering appeared to consistently reach this maximum speed. However, later, after upgrading to the Meta Quest 3, which has a higher resolution of 2064x2208 pixels per eye, we noticed the framerate did drop below 72 Hz, motivating the foveated rendering feature described here and a formal evaluation of rendering performance.

Foveated rendering strives to maintain a high quality rendering in the focal region of the eye(s) and decrease rendering quality in the periphery. It has been employed previously for volume rendering (e.g., [1], [8], [40]) and is a common technique for VR rendering both with and without eye tracking [29], [45]. Our approach does not use eye tracking. We use a 35-degree focal area, meaning that for eye, all pixels within 17.5 degrees of center of the filmplane are rendered at full quality. Then, for each pixel beyond this circle, the rendering quality is decreased smoothly by increasing the step size used in the volume rendering ray traversal in proportion to the pixel's distance from the focal region. The effect is seen in the middle row of Fig. 12–note the subtle difference in the small inset images, which zoom in on pixels in the periphery.

To evaluate the rendering performance, we ran a systematic test with the same data and visualization settings as in Fig. 11. We used the default settings for the Meta Quest Link software with the exception of disabling Asynchronous SpaceWarp and collected framerate data in repeated trials while varying two parameters: (1) number of plans displayed simultaneously, which we varied, just as in the mock plan review meeting, from 1 to 4, and (2) scale of the plans, which we varied from 1 (matching real-world size) to 4 (a reasonable maximum for comfortable viewing). Since the framerate can vary as the head moves and new volumes come into view, we moved the headset consistently during each trial by placing the headset on a tripod that is rotated by hand to focus on one plan for 5 seconds, then rotate to focus on the next plan, and so on for all of the plans displayed in trial moving from left-to-right, then repeating the pattern moving from right-to-left. To avoid frame-to-frame noise, we calculated the framerate as the average of 1.0/Time.deltaTime (as reported by Unity each frame) within consecutive 0.5 second intervals and write these averages to a file. This yielded hundreds of framerate samples per trial; thus, we report per-trial medians, interquartile ranges, and outliers in the box-and-whisker plots in Fig. 12.

Using the foveated rendering technique (plotted in blue in Fig. 12), the framerate consistently reaches 72 Hz (the default max framerate for Meta Quest Link), only dips below for data points designated as outliers, and never decreases below 60 Hz. Without foveated rendering (plotted in red), the framerate reaches 72 Hz when visualizing up to four volumes at a scale of 1.0, but it begins to drop when two volumes are visualized at a scale of 3.0, three volumes at a scale of 2.0, or four volumes at a scale of 1.5. In the extreme cases, the framerate can drop by 50% or more. Our data show a lower cap at 36 Hz, but we believe this is a nuance of working with the proprietary Meta Quest Link software—we assume it is dropping frames in these situations and the true framerate is lower.

The performance of the rendering techniques introduced in this paper is closely correlated with the resolution of the displays as well as the number of volumes rendered simultaneously and the scale of the volumes. Also, the techniques combine well with foveated rendering as one way to increase performance to support higher resolution displays.

VII. Future Work and Conclusion

While frame rate has been reported, we acknowledge that latency data has not been measured. Two kinds of latency data could be relevant for future study: tracking and networking. With tracking latency, a dedicated camera setup could be developed to record real-world hand movements and compare those with their virtual counterparts. For networking latency, to date, we have only deployed the system on local networks, but the potential to do remote collaborative treatment planning is exciting. All data are cached locally and all rendering is performed locally; so, these aspects should translate naturally. Separating the clients by distance would add latency to the network-based synchronization of the visualization state, and this may require extensions to the user interface to maintain effective collaborative workflows.

The most impactful future work may come from scaling the comparative visualization up to tens or hundreds of plans. The

Authorized licensed use limited to: University Of Minnesota Duluth. Downloaded on June 23,2026 at 16:38:47 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Visualization and Computer Graphics. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2026.3686821

clinical team has developed multicriteria optimization tools for generating and working with such ensemble datasets, and we envision the immersive visualization could be linked with complementary views of the ensemble and machine learning that guides users toward the most interesting subsets of plans to compare through immersive visualization.

We conclude that in the context of this radiation treatment planning center at the Mayo Clinic, the collaborative immersive visualization system described in this paper shows promise as a powerful tool for dosimetrists, medical physicists, and physicians to analyze treatment plans. The expected benefits relative to the team's current practice include noticing patterns more efficiently (e.g., locating hot spots), understanding the "big picture" (i.e., how plans impact the many structures within a treatment volume), and facilitating discussion of spatial data features. Despite the application-driven and context-dependent nature of this design study, we can synthesize several guidelines for visualization and VR researchers and practitioners. Adding to the scientific and theoretical discussion of how to most effectively design interactive 3D visualization systems like this, we provide the following guidelines: 1. Combine smart cursors and easily understandable widget-based user interfaces with bimanual and gestural user interfaces that are less self-revealing but are fast for expert users. 2. Combine 3D visualization with linked complementary 2D visualizations (e.g., DVH plots, slicing planes). 3. Combine GPU-accelerated volume and surface rendering techniques to help users accurately perceive dose distributions within and near anatomical structures. 4. Support visual comparison through juxtaposition, interchangeable views that animate over time, explicit visual encodings of calculated differences, or, ideally, combinations of these approaches. The work also extends well-established perception-enhancing spatial visualization styles (e.g., curvature following lines) by introducing modern GPU-based implementations that enable interactive rendering with dynamic updates and superimposed volume rendering.

ACKNOWLEDGEMENTS

The authors thank physicians, David Routman, MD; William Breen, MD; Ted Hoene; and the whole Mayo Clinic Radiation Treatment team for valuable feedback.

REFERENCES

[1] T. Ananpiriyakul, J. Anghel, K. Potter, A. Joshi, J. Anghel, K. Potter, and A. Joshi. A Gaze-Contingent System for Foveated Multiresolution Visualization of Vector and Volumetric Data. *Electronic Imaging*, 32:1–11, 2020. doi: 10.2352/ISSN.2470-1173.2020.1.VDA-374 11

[2] A. S. Bair, D. H. House, and C. Ware. Texturing of layered surfaces for optimal viewing. *IEEE transactions on visualization and computer graphics*, 12(5):1125–1132, 2006. doi: 10.1109/TVCG.2006.183 3

[3] H. Bannister, B. Selwyn-Smith, C. Anslow, B. Robinson, P. Kane, and A. Leong. Collaborative VR Simulation for Radiation Therapy Education. In J.-F. Uhl, J. Jorge, D. S. Lopes, and P. F. Campos, eds., *Digital Anatomy : Applications of Virtual, Mixed and Augmented Reality*, Human–Computer Interaction Series, pp. 199–221. Springer International Publishing, Cham, 2021. doi: 10.1007/978-3-030-61905-3_11 3

[4] S. Benford, J. Bowers, L. E. Fahlén, C. Greenhalgh, and D. Snowdon. User embodiment in collaborative virtual environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '95*, pp. 242–249. ACM Press, Denver, Colorado, United States, 1995. doi: 10.1145/223904.223935 3

[5] A. Boejen and C. Grau. Virtual reality in radiation therapy training. *Surgical oncology*, 20(3):185–188, 2011. 3

[6] J. Bowers, R. Wang, L.-Y. Wei, and D. Maletz. Parallel poisson disk sampling with spectrum analysis on surfaces. *ACM Transactions on Graphics (TOG)*, 29(6):1–10, 2010. doi: 10.1145/1882261.1866188 7

[7] B. Brown, S. Reeves, and S. Sherwood. Into the wild: challenges and opportunities for field trial methods. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pp. 1657–1666, 2011. 8

[8] V. Bruder, C. Schulz, R. Bauer, S. Frey, D. Weiskopf, and T. Ertl. Voronoi-Based Foveated Volume Rendering. *EuroVis 2019 - Short Papers*, 2019. doi: 10.2312/evs.20191172 11

[9] A. K. Bryant, M. P. Banegas, M. E. Martinez, L. K. Mell, and J. D. Murphy. Trends in radiation therapy among cancer survivors in the united states, 2000–2030. *Cancer Epidemiology, Biomarkers & Prevention*, 26(6):963–970, 2017. 1

[10] S. Carpendale. Evaluating Information Visualizations. In *Information Visualization*. Springer Berlin Heidelberg, 2008. doi: 10.1007/978-3-540-70956-5_2 3

[11] T. Chandler, M. Cordeil, T. Czauderna, T. Dwyer, J. Glowacki, C. Goncu, M. Klapperstueck, K. Klein, K. Marriott, F. Schreiber, and E. Wilson. Immersive Analytics. In *2015 Big Data Visual Analytics (BDVA)*, pp. 1–8. IEEE, Hobart, Australia, Sept. 2015. doi: 10.1109/BDVA.2015.7314296 3

[12] V. Chheang, P. Saalfeld, T. Huber, F. Huettl, W. Kneist, B. Preim, and C. Hansen. Collaborative Virtual Reality for Laparoscopic Liver Surgery Training. In *2019 IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR)*, pp. 1–17. IEEE, San Diego, CA, USA, Dec. 2019. doi: 10.1109/AIVR46125.2019.00011 3

[13] V. K. Chowdhry, Z. Daniel, and R. Hackett. Departmental quality initiative to establish turnaround times from simulation to treatment. *Technical Innovations & Patient Support in Radiation Oncology*, 19:37–40, 2021. 1

[14] J. C. Chu, X. Gong, Y. Cai, M. C. Kirk, T. W. Zusag, S. Shott, M. J. Rivard, C. S. Melhus, G. A. Cardarelli, A. Hurley, et al. Application of holographic display in radiotherapy treatment planning ii: a multi-institutional study. *Journal of Applied Clinical Medical Physics*, 10(3):115–124, 2009. 3

[15] B. D. Conner, S. S. Snibbe, K. P. Herndon, D. C. Robbins, R. C. Zeleznik, and A. Van Dam. Three-dimensional widgets. In *Proceedings of the 1992 Symposium on Interactive 3D Graphics - SI3D '92*, pp. 183–188. ACM Press, Cambridge, Massachusetts, United States, 1992. doi: 10.1145/147156.147199 3

[16] F. Cosentino, N. W. John, and J. Vaarkamp. Rad-ar: Radiotherapy - augmented reality. In *2017 International Conference on Cyberworlds (CW)*, pp. 226–228, 2017. doi: 10.1109/CW.2017.56 3

[17] S. Glaser, B. Warfel, J. Price, J. Sinacore, and K. Albuquerque. Effectiveness of Virtual Reality Simulation Software in Radiotherapy Treatment Planning Involving Non-Coplanar Beams with Partial Breast Irradiation as a Model. *Technology in Cancer Research & Treatment*, 11(5):409–414, Oct. 2012. doi: 10.7785/tcrt.2012.500256 3

[18] M. Gleicher. Considerations for Visualizing Comparison. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):413–423, Jan. 2018. Conference Name: IEEE Transactions on Visualization and Computer Graphics. doi: 10.1109/TVCG.2017.2744199 3

[19] M. Gleicher, D. Albers, R. Walker, I. Jusufi, C. D. Hansen, and J. C. Roberts. Visual comparison for information visualization. *Information Visualization*, 10(4):289–309, Oct. 2011. Publisher: SAGE Publications. doi: 10.1177/1473871611416549 2, 3

[20] T. Grossman and R. Balakrishnan. The bubble cursor: Enhancing target acquisition by dynamic resizing of the cursor's activation area. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '05, pp. 281–290. Association for Computing Machinery, New York, NY, USA, Apr. 2005. doi: 10.1145/1054972.1055012 4

[21] G. D. Haan, M. Koutek, and F. H. Post. IntenSelect: Using Dynamic Object Rating for Assisting 3D Object Selection. In *Eurographics Symposium on Virtual Environments*, p. 9 pages. [object Object], 2005. doi: 10.2312/EGVE/IPT_EGVE2005/201-209 3

[22] D. E. Holt, A. L. Carr, S. Roberts, S. A. Milgrom, E. Kolva, B. D. Kavanagh, G. E. Switzer, C. Eitel, J. Nelson, B. Miller, J. Shiao, A. C. Mueller, S. D. Karam, and T. R. Clapp. The use of virtual reality for 3d diagnostic imaging review enhances understanding and education of patients with cancer. *International Journal of Radiation Oncology*Biology*Physics*, 122(5):1102–1112, 2025. doi: 10.1016/j.ijrobp.2025.02.038 3

[23] R. Homan. SmartScene: An Immersive, Realtime, Assembly, Verification and Training Application. In *Computational Tools and Facilities*

Authorized licensed use limited to: University Of Minnesota Duluth. Downloaded on June 23,2026 at 16:38:47 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Visualization and Computer Graphics. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2026.3686821

for the Next-Generation Analysis and Design Environment, pp. 119–140, Mar. 1997. 4

[24] V. Interrante. Illustrating surface shape in volume data via principal direction-driven 3D line integral convolution. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques - SIGGRAPH '97*, pp. 109–116. ACM Press, Not Known, 1997. doi: 10.1145/258734.258796 7

[25] V. Interrante, H. Fuchs, and S. Pizer. Illustrating transparent surfaces with curvature-directed strokes. In *Proceedings of Seventh Annual IEEE Visualization '96*, pp. 211–218, Oct. 1996. doi: 10.1109/VISUAL.1996.568110 2, 3

[26] V. Interrante, H. Fuchs, and S. Pizer. Conveying the 3D shape of smoothly curving transparent surfaces via texture. *IEEE Transactions on Visualization and Computer Graphics*, 3(2):98–117, June 1997. doi: 10.1109/2945.597794 2, 3

[27] V. Interrante, H. Fuchs, and S. Pizer. Conveying the 3D shape of smoothly curving transparent surfaces via texture. *IEEE Transactions on Visualization and Computer Graphics*, 3(2):98–117, April-June/1997. doi: 10.1109/2945.597794 7

[28] P. Isenberg, N. Elmqvist, J. Scholtz, D. Cernea, K.-L. Ma, and H. Hagen. Collaborative visualization: Definition, challenges, and research agenda. *Information Visualization*, 10(4):310–326, Oct. 2011. doi: 10.1177/1473871611412817 3

[29] S. Jabbireddy, X. Sun, X. Meng, and A. Varshney. Foveated Rendering: Motivation, Taxonomy, and Research Directions, Nov. 2022. doi: 10.48550/arXiv.2205.04529 11

[30] P. Khanal, A. Vankipuram, A. Ashby, M. Vankipuram, A. Gupta, D. Drumm-Gurnee, K. Josey, L. Tinker, and M. Smith. Collaborative virtual reality based advanced cardiac life support training simulator using virtual reality principles. *Journal of Biomedical Informatics*, 51:49–59, Oct. 2014. doi: 10.1016/j.jbi.2014.04.005 2, 3

[31] K. Kim, J. V. Carlis, and D. F. Keefe. Comparison techniques utilized in spatial 3D and 4D data visualizations: A survey and future directions. *Computers & Graphics*, 67:138–147, Oct. 2017. doi: 10.1016/j.cag.2017.05.005 2, 3

[32] K. Kim, B. Jackson, I. Karamouzas, M. Adeagbo, S. J. Guy, R. Graff, and D. F. Keefe. Bema: A multimodal interface for expert experiential analysis of political assemblies at the Pnyx in ancient Greece. In *2015 IEEE Symposium on 3D User Interfaces (3DUI)*, pp. 19–26, Mar. 2015. doi: 10.1109/3DUI.2015.7131720 3

[33] A. Kunert, A. Kulik, S. Beck, and B. Froehlich. Photoportals: shared references in space and time. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*, pp. 1388–1399. ACM, Baltimore Maryland USA, Feb. 2014. doi: 10.1145/2531602.2531727 2, 3

[34] A. Kunert, T. Weissker, B. Froehlich, and A. Kulik. Multi-Window 3D Interaction for Collaborative Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics*, 26(11):3271–3284, Nov. 2020. doi: 10.1109/TVCG.2019.2914677 3

[35] G. Kurtenbach and William Buxton. User learning and performance with marking menus. *Proceedings of the SIGCHI conference on Human factors in computing systems.*, pp. 258–264, Apr. 1994. 3

[36] B. M. Kyaw, N. Saxena, P. Posadzki, J. Vseteckova, C. K. Nikolaou, P. P. George, U. Divakar, I. Masiello, A. A. Kononowicz, N. Zary, and L. Tudor Car. Virtual Reality for Health Professions Education: Systematic Review and Meta-Analysis by the Digital Health Education Collaboration. *Journal of Medical Internet Research*, 21(1):e12959, Jan. 2019. doi: 10.2196/12959 3

[37] J. J. LaViola Jr., E. Kruijff, R. P. McMahan, D. Bowman, and I. P. Poupyrev. *3D User Interfaces: Theory and Practice*. Addison-Wesley Professional, Apr. 2017. 3, 4

[38] B. Lee, X. Hu, M. Cordeil, A. Prouzeau, B. Jenny, and T. Dwyer. Shared Surfaces and Spaces: Collaborative Data Visualisation in a Co-located Immersive Environment. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1171–1181, Feb. 2021. arXiv:2009.00050 [cs]. doi: 10.1109/TVCG.2020.3030450 3

[39] M. Levoy, H. Fuchs, S. Pizer, J. Rosenman, E. Chaney, G. Sherouse, V. Interrante, and J. Kiel. Volume rendering in radiation treatment planning. In *[1990] Proceedings of the First Conference on Visualization in Biomedical Computing*, pp. 4–10. IEEE Comput. Soc. Press, Atlanta, GA, USA, 1990. doi: 10.1109/VBC.1990.109295 2

[40] M. Levoy and R. Whitaker. Gaze-directed volume rendering. In *Proceedings of the 1990 Symposium on Interactive 3D Graphics*, I3D '90, pp. 217–223. Association for Computing Machinery, 1990. doi: 10.1145/91385.91449 11

[41] C. Lundström, P. Ljung, A. Persson, and A. Ynnerman. Uncertainty visualization in medical volume rendering using probabilistic animation. *IEEE transactions on visualization and computer graphics*, 13(6):1648–1655, 2007. doi: 10.1109/TVCG.2007.70518 3

[42] S. LYi, J. Jo, and J. Seo. Comparative Layouts Revisited: Design Space, Guidelines, and Future Directions. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1525–1535, Feb. 2021. Conference Name: IEEE Transactions on Visualization and Computer Graphics. doi: 10.1109/TVCG.2020.3030419 2, 3

[43] D. P. Mapes and J. M. Moshell. A Two-Handed Interface for Object Manipulation in Virtual Environments. *Presence: Teleoperators and Virtual Environments*, 4(4):403–416, Nov. 1995. doi: 10.1162/pres.1995.4.4.403 3

[44] D. P. Mapes and J. M. Moshell. A Two-Handed Interface for Object Manipulation in Virtual Environments. *Presence: Teleoperators and Virtual Environments*, 4(4):403–416, Nov. 1995. doi: 10.1162/pres.1995.4.4.403 4

[45] B. Mohanto, A. T. Islam, E. Gobbetti, and O. Staadt. An integrative view of foveated rendering. *Computers & Graphics*, 102:474–501, 2022. doi: 10.1016/j.cag.2021.10.010 11

[46] J. W. Nam, K. McCullough, J. Tveite, M. M. Espinosa, C. H. Perry, B. T. Wilson, and D. F. Keefe. Worlds-in-Wedges: Combining Worlds-in-Miniature and Portals to Support Comparative Immersive Visualization of Forestry Data. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 747–755. IEEE, Osaka, Japan, Mar. 2019. doi: 10.1109/VR.2019.8797871 3

[47] C. North. Toward measuring visualization insight. *IEEE Computer Graphics and Applications*, 26(3):6–9, May 2006. doi: 10.1109/MCG.2006.70 9

[48] D. Patel, L. P. Muren, A. Mehus, Y. Kvinnsland, D. M. Ulvang, and K. P. Villanger. A virtual reality solution for evaluation of radiotherapy plans. *Radiotherapy and Oncology*, 82(2):218–221, Feb. 2007. doi: 10.1016/j.radonc.2006.11.024 3

[49] I. Poupyrev, M. Billinghurst, S. Weghorst, and T. Ichikawa. The go-go interaction technique: Non-linear mapping for direct manipulation in VR. In *Proceedings of the 9th Annual ACM Symposium on User Interface Software and Technology*, UIST '96, pp. 79–80. Association for Computing Machinery, New York, NY, USA, Nov. 1996. doi: 10.1145/237091.237102 3

[50] B. Preim, R. Raidou, N. Smit, and K. Lawonn. Chapter 8 - visual computing for radiation treatment planning. In B. Preim, R. Raidou, N. Smit, and K. Lawonn, eds., *Visualization, Visual Analytics and Virtual Reality in Medicine*, The MICCAI Society Book Series, pp. 199–221. Academic Press, 2023. doi: 10.1016/B978-0-12-822962-0.00016-X 3

[51] K. Riege, T. Holtkamper, G. Wesche, and B. Frohlich. The Bent Pick Ray: An Extended Pointing Technique for Multi-User Interaction. In *3D User Interfaces (3DUI'06)*, pp. 62–65, Mar. 2006. doi: 10.1109/VR.2006.127 4

[52] G. Rong and T.-S. Tan. Jump flooding in GPU with applications to Voronoi diagram and distance transform. In *Proceedings of the 2006 Symposium on Interactive 3D Graphics and Games*, I3D '06, pp. 109–116. Association for Computing Machinery, New York, NY, USA, Mar. 2006. doi: 10.1145/1111411.1111431 6

[53] P. Saraiya, C. North, and K. Duca. An insight-based methodology for evaluating bioinformatics visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 11(4):443–456, July 2005. doi: 10.1109/TVCG.2005.53 9

[54] M. Schlachter, R. Raidou, L. Muren, B. Preim, P. Putora, and K. Bühler. State-of-the-Art Report: Visual Computing in Radiation Therapy Planning. *Computer Graphics Forum*, 38(3):753–779, 2019. doi: 10.1111/cgf.13726 1, 3

[55] M. Sedlmair, M. Meyer, and T. Munzner. Design study methodology: Reflections from the trenches and the stacks. *IEEE transactions on visualization and computer graphics*, 18(12):2431–2440, 2012. 1, 2, 7

[56] R. L. Siegel, K. D. Miller, H. E. Fuchs, and A. Jemal. Cancer Statistics, 2021. *CA: a cancer journal for clinicians*, 71(1):7–33, Jan. 2021. doi: 10.3322/caac.21654 1

[57] R. Skarbez, F. P. Brooks, Jr., and M. C. Whitton. A Survey of Presence and Related Concepts. *ACM Computing Surveys*, 50(6):96:1–96:39, Nov. 2017. doi: 10.1145/3134301 3

[58] R. Skarbez, N. F. Polys, J. T. Ogle, C. North, and D. A. Bowman. Immersive Analytics: Theory and Research Agenda. *Frontiers in Robotics and AI*, 6:82, Sept. 2019. doi: 10.3389/frobt.2019.00082 3

[59] T. Su, W. Sung, C. Jiang, S. Sun, and C. Wu. The Development of a VR-Based Treatment Planning System for Oncology. In *2005 IEEE Engineering in Medicine and Biology 27th Annual Conference*, pp. 6104–6107. IEEE, Shanghai, China, 2005. doi: 10.1109/IEMBS.2005.1615886 3

Authorized licensed use limited to: University Of Minnesota Duluth. Downloaded on June 23,2026 at 16:38:47 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Visualization and Computer Graphics. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2026.3686821

[60] A. Tang, M. Tory, B. Po, P. Neumann, and S. Carpendale. Collaborative coupling over tabletop displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '06, pp. 1181–1190. Association for Computing Machinery, New York, NY, USA, Apr. 2006. doi: 10.1145/1124772.1124950 3

[61] C.-Y. Wang, T.-F. Lee, and C.-H. Fang. A volume visualization system with augmented reality interaction for evaluation of radiotherapy plans. In *2009 Fourth International Conference on Innovative Computing, Information and Control (ICICIC)*, pp. 433–436. IEEE, 2009. 3

[62] L. J. Wang, B. Casto, J. Y. Luh, and S. J. Wang. Virtual reality-based education for patients undergoing radiation therapy. *Journal of Cancer Education*, 37(3):694–700, 2022. 3

[63] J. W. Ward, R. Phillips, A. Boejen, C. Grau, D. Jois, and A. W. Beavis. A Virtual Environment for Radiotherapy Training and Education - VERT. *Eurographics*, 2011. doi: 10.2312/EG2011/med/005-008 3

[64] H. Yang and G. M. Olson. Exploring collaborative navigation: The effect of perspectives on group performance. In *Proceedings of the 4th International Conference on Collaborative Virtual Environments*, CVE '02, pp. 135–142. Association for Computing Machinery, New York, NY, USA, Sept. 2002. doi: 10.1145/571878.571899 3

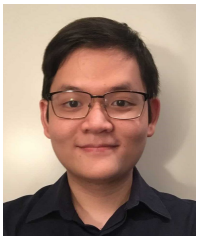

Kiet Tran is a Ph.D. student in Computer Science at the University of Minnesota, Twin Cities. His research interests include mixed-reality tools for creative processes, scientific visualization, and data storytelling. He is particularly interested in designing interfaces for 3D drawing in mixed reality. He received his B.A. in Computer Science from Macalester College.

Matthias Broske is an AR/VR software engineer at Apple. He received a B.S. in Computer Science from the University of Minnesota, Twin Cities, where he worked as an undergraduate researcher in the Interactive Visualization Lab. His research interests include scientific visualization, GPU-accelerated rendering, and interactive computer graphics.

Michael G. Herman is a Professor of Medical Physics at the Mayo Clinic. He received his Doctor of Philosophy at the University of Rochester. He is board-certified at Therapeutic Medical Physics. His research interests are medical physics, radiation oncology, and stereotactic radiosurgery.

Victoria Interrante is a Full Professor in the Department of Computer Science and Engineering at the University of Minnesota, Twin Cities. Her current research centers on improving the human experience in immersive virtual environments, including efforts related to mitigating cybersickness, addressing racial bias, and harnessing the potential of virtual natural environments to improve physical and mental health and well-being. In 2020, she was honored with the IEEE VGTC Virtual Reality Career Award for her lifetime contributions to visualization and visual perception for augmented and virtual reality, and she was inducted in 2022 into the IEEE VGTC Virtual Reality Academy.

Evan Suma Rosenberg is an Associate Professor in the Department of Computer Science and Engineering at the University of Minnesota. Before University of Minnesota, they completed postdoctoral work at the University of Southern California and went on to serve as a senior research associate and research assistant professor at the university. Suma Rosenberg earned the George W. Taylor Career Development Award in 2022.

Daniel W. Mundy is a Medical Physicist and Assistant Professor in the Department of Radiation Oncology at Mayo Clinic in Rochester, Minnesota. He is currently the Clinical Practice Chair of the Radiation Oncology Physics Division, Co-chair of the Radiation Oncology IT Group, Assistant Director of the Proton Medical Physics fellowship program, and member of the institutional Medical Device Oversight Subcommittee. His research focuses on the practical aspects of proton treatment planning and delivery. Prior to joining the Mayo Clinic staff, Dr. Mundy graduated with B.S. and M.S. degrees in Nuclear Engineering from Purdue University, earned a Ph.D. in Biomedical Engineering from the Mayo School of Graduate Medical Education, and completed a clinical Medical Physics Fellowship at Mayo Clinic.

Daniel F. Keefe is a Distinguished University Teaching Professor and Professor of Computer Science and Engineering at the University of Minnesota, Twin Cities. His research centers on scientific data visualization and interactive computer graphics. In addition to his work in computer science, Keefe is also an accomplished artist and has published and exhibited work in top international venues for digital art. Before joining the University of Minnesota, Keefe did post-doctoral work at Brown University jointly with the departments of Computer Science and Ecology and Evolutionary Biology and with the Rhode Island School of Design. He received the Ph.D. in 2007 from Brown University's Department of Computer Science and the B.S. in Computer Engineering summa cum laude from Tufts University in 1999.

Authorized licensed use limited to: University Of Minnesota Duluth. Downloaded on June 23,2026 at 16:38:47 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.