

VolumeVisual: Design and Evaluation of an Educational Software Tool for Teaching and Learning Volume Visualization

Xueyi Bao, Jun Han, Chaoli Wang
University of Notre Dame

Abstract

Visualizing three-dimensional volumetric data sets, often referred to as volume visualization, has been an essential topic in scientific visualization. Volume visualization is concerned with visualizing scalar field data (e.g., temperature, pressure, intensity). Existing software tools can perform interactive visualization of volumetric data sets. However, these tools are often sophisticated and mainly geared towards researchers and developers. In this paper, we present VolumeVisual, a web-based educational software tool for teaching and learning volume visualization. VolumeVisual integrates the two most popular volume visualization techniques, direct volume rendering and isosurface rendering, to give students a comprehensive understanding of basic concepts and algorithms. The software tool includes a study component and a quiz component. Effectiveness of VolumeVisual is evaluated by conducting a formal user study consisting of an introduction and training session, a survey, followed by a quiz.

1 Introduction

As visualization has become an indispensable means for analyzing data generated from applications that span various STEM fields, more research efforts pay attention to education in data science and human-centered computing. Following this trend is a significant need for high-quality educational tools for teaching and learning data visualization. Over the years, different tools have been developed for information visualization (InfoVis) that span across software framework [8], domain-specific language [5], and visualization education [20, 10]. However, there is a lack of educational tools for teaching and learning scientific visualization (SciVis). Unlike InfoVis, which handles non-spatial data like node-link diagrams, SciVis deals with 3D spatial data that needs more complex visualization and rendering techniques stemming from computer graphics. One important area in SciVis is volume visualization (VolVis), of which primary techniques, such as direct volume rendering (DVR) and isosurface rendering (IR), are nontrivial to implement and challenging to learn. Usually, these techniques are taught by instructors through orally explaining concepts and definitions, drawing diagrams and illustrations, and playing custom-made animations or video clips. Interactive VolVis examples running on an educational tool can facilitate this teaching and learning process to a large extent. Therefore, both instructors and students can potentially benefit from using an interactive pedagogical tool for VolVis.

In this paper, we present VolumeVisual, a web-based educational software tool that illustrates essential concepts and algorithms in VolVis. DVR and IR are utilized to enrich students' understanding of volumetric data sets generated from real-world medical applications and scientific simulations. Students can evaluate and compare these two different techniques in side-by-side panels to better grasp the pros and cons of each method.

The motivation of VolumeVisual derives from two different aspects of visualization education. The first aspect is to assist students in learning sophisticated VolVis techniques efficiently. Many existing packages, such as VTK and ParaView, encompass diverse VolVis techniques, which takes a significant effort for beginners to grasp necessary intricacies in order to use the tool and produce meaningful results. This unavoidably deviates the learning focus from understanding visualization concepts to mastering a complex package. Our design mainly targets beginning learners, and we aim to provide a clean yet straightforward interface without distracting visuals and irrelevant details. Students will only focus on learning the key concepts and basic algorithms in VolVis, which will facilitate their understanding. The second aspect is to teach students to use VolVis techniques to solve practical problems. VolumeVisual provides

students with much flexibility to interact with the tool and explore the visualization. For example, students can learn how to tweak the transfer function to generate a desirable visualization and how to crop the volume to reveal interesting internal volumetric features. The effectiveness of VolumeVisual was evaluated via a user study, including a self-study, survey, and quiz.

2 Related Work

VolVis has been around for decades and remains an active area of research in SciVis. It is essential to scientific, engineering, and medical applications that require visualization of 3D volumetric data sets. Examples include visualization of data acquired from medical imaging devices or produced from scientific simulations. In VolVis, a single isosurface based on a given isovalue can be extracted and rendered through IR, or the entire volume can be rendered through DVR without constructing any intermediate geometry.

For IR, VolumeVisual utilizes the marching cube algorithm [15], the most common method for isosurface extraction and leverages the GPU-accelerated isosurface similarity map method [11] to identify representative isosurfaces. The triangle meshes extracted from the 3D scalar field are then rendered with traditional geometric rendering techniques. Isosurface similarity map was first introduced in [6]. It is a symmetric matrix recording the similarity values for each pair of sampled isosurfaces.

Unlike IR, DVR generates images of a 3D volumetric data set without explicitly extracting geometry from the data. Nowadays, most of the DVR techniques fall into two categories: texture-based DVR [24, 12, 14] and raycasting DVR [13]. VolumeVisual implements a GPU-based ray casting solution [13]. In raycasting DVR, the ray traverses through the 3D scalar field, and intensity is sampled along the ray inside the volume. Combined with simple illumination, raycasting DVR can create high-quality volumetric rendering effects.

Although there is a widespread and unchallenged belief that visualization is useful in both teaching and learning [21], we still lack pedagogical tools for teaching visualization. Back in 2005, “Education and Training” was listed as one of the top ten unsolved problems in InfoVis [7]. The overarching theme for the 2015 Gordon Research Conference on Visualization in Science and Education is “Grand Challenges in the Use of Visualization in Science and Education” [1]. More recently, educational psychology concepts were applied to design four types of online guides for InfoVis [20]. However, few works have been done for SciVis. Hertzberg and Sweetman [9] designed a flow visualization course focusing on studio/laboratory experiences. Wang et al. [23] presented an education tool named FlowVisual for teaching and learning 2D flow visualization and later extended it to 3D flow visualization [22]. Work focusing on VolVis education is even more scarce. There are some existing tools such as VTK [2] and ParaView [3] for VolVis. However, these are comprehensive visualization packages for advanced users and provide abundant visualization techniques that may quickly overwhelm beginning students. They could distract students from grasping important visualization concepts efficiently. To minimize students’ learning difficulties, VolumeVisual only integrates the most important algorithms and functions in VolVis and is specifically designed for college students in the general area of data science to learn VolVis effectively.

3 Design Principles

Novel visualization tools should be based on the theoretical tenets of *constructivism* and *constructionism* [19]. The approach of constructivism advocates that students learn best when given opportunities to tie together their knowledge with classroom material to discover or deduce desired principles for themselves [18]. The approach of constructionism posits that students are better able to apprehend complicated concepts through the act of building their own models of the concepts with the aid of computer-based technologies [17]. Furthermore, cognitive load theory [16] describes the use of structures for acquiring, retrieving, and processing information with specific emphasis on long-term and working memories. Based on the above learning theory, we designed VolumeVisual with the following principles.

P1: Simplicity. The goal of VolumeVisual is to explain key concepts and illustrate basic algorithms in VolVis to facilitate understanding. Unlike any other complex visualization packages, VolumeVisual does not include a rich set of visualization algorithms and techniques which may overwhelm students. VolumeVisual only implements DVR and IR, the two most representative VolVis techniques. It provides a simple and easy-to-use interface so that users can

indeed gain insights quickly. Doing so limits unnecessary or even distracting information and hence, helps learners focus on the most salient features of the functions.

P2: Interactivity. Interactivity is the key to any visualization tool. VolumeVisual provides students with the opportunity to generate knowledge in an iterative fashion as they become actively involved in the process of interactive exploration. Interactivity also directly relates to the notion of cognitive activity. As a visualization toolkit, VolumeVisual goes beyond what animations can offer by providing the capability for users to interact with the software tool. Students can change parameters for lighting and viewing, rotate the view, and tweak the transfer function for DVR, which builds a highly interactive environment where knowledge could be constructed by students more efficiently based on the constructivism learning theory. Furthermore, default lighting parameters will be used to generate desired rendering results so that users can focus on the interaction with respect to surface and volume mapping.

P3: Comparison. VolumeVisual supports simultaneous visualization of two different VolVis techniques applied to the same volumetric data set in two side-by-side display panels, allowing students to make an easy comparison by adjusting parameters such as cutting planes, lighting, and viewing. The two display panels also react synchronically to interactions when parameters are changed. This is more helpful for students to understand the respective strengths and weaknesses of these rendering techniques rather than solely relying on instructors in explaining the differences.

P4: Modular and web-friendly. We design VolumeVisual with possible future expansion in mind. Therefore, we follow a modular design, making it easy to include other additions not yet planned. WebGL is used for easy deployment, which is based on the widely-adopted Internet standard of HTML5. Content developed under the new standards can run on all major browsers without the need for specialized plugins and can be automatically adapted to viewing on a variety of mobile devices.

4 Overview

We design VolumeVisual to engage students in active learning of the core concepts and algorithms of VolVis techniques. Users can explore VolVis results from various data sets generated from medical applications and scientific simulations. VolumeVisual provides two side-by-side panels for DVR and IR, combined with different interactive functions to enable students to understand, compare, and evaluate these two important VolVis algorithms. The primary interactions in VolumeVisual include switching between data sets, changing different rendering parameters (lighting, viewing, etc.), and specifying the transfer function. VolumeVisual has two components, the *study component* and *quiz component*. VolumeVisual is implemented using HTML, CSS, and JavaScript, WebGL, and D3.js.




Figure 1: The user interface of VolumeVisual (study component). (a) Screenshot with Cutting Plane unfolded on the control panel. (b) Unfolded Lighting Parameters and Viewing Parameters controls. (c) Unfolded Volume Rendering and Isosurface Rendering controls along with the display of data set information.

Figure 1 (a) shows a screenshot of the VolumeVisual interface with three major parts: a header bar on the top, two side-by-side rendering panels, and a scrollable control panel on the right. The top bar displays the VolumeVisual logo and hints about how to change the view with mouse, keyboard, or touchpad. The two side-by-side rendering panels show DVR and IR results, respectively, with the transfer function editor under the DVR panel and the color picker under the IR panel. At the top of the control panel, a dropdown list enables data selection. Two links **about** and **help** are for reviewing the “about” page information and the tool’s onboarding information, respectively. Beneath, there are main control panel sections, including five groups of control (from top to bottom, Cutting Plane, Lighting Parameters, Viewing Parameters, Volume Rendering, and Isosurface Rendering), followed by the description of the currently chosen data set. These groups of control can be folded or unfolded as needed. Besides, users can click on the **Reset All** button to reset all the parameters to their default settings. Table 1 lists the data sets used in the study component (left) and quiz component (right).

Data Set (Variable)	Domain	Data Set (Variable)	Domain
Argon Bubble	Scientific	Atmospheric	Scientific
Buckyball	Scientific	Bonsai	Medical
Climate (SALT)	Scientific	Combustion (HR)	Scientific
Climate (TEMP)	Scientific	Combustion (OH)	Medical
Frog	Medical	Five Jets	Scientific
Ionization (H)	Scientific	Foot	Scientific
Ionization (PD)	Scientific	Hurricane Isabel (VAPOR)	Scientific
Northridge Earthquake	Scientific	Hurricane Isabel (WSMAG)	Scientific
Skull	Medical	Lobster	Medical
Solar Plume	Scientific	Tornado	Scientific
Supernova	Scientific	UNC Brain	Medical
Visible Woman	Medical		
Vortex	Scientific		

Table 1: The data sets included in VolumeVisual. The left and right tables show the data sets used in the study and quiz components, respectively. Four data sets (Climate, Combustion, Hurricane Isabel, and Ionization) are multivariate data and the corresponding variables used are listed.

4.1 IR and DVR

As VolumeVisual incorporates both IR and DVR for VolVis. In the following, we briefly introduce these two techniques and their implementation details.

IR. Similar to an *isocontour*, which is a contour in 2D space that intersects pixels having a constant value, an *isosurface* is a surface in 3D space that intersects voxels having a constant value. To allow interactive IR at runtime, VolumeVisual performs isosurface extraction offline as a preprocessing. Given a volumetric data set, we extract a set of isosurfaces at uniformly chosen isovalues, from which we identify a small subset as the representatives [11]. An isosurface only shows the information with respect to the given isovalue. In VolumeVisual, using a slider, users can go through each of the representative isosurfaces to get a comprehensive understanding of the underlying volumetric data set. A simple Blinn-Phong illumination model is used to render the output triangle primitives of isosurfaces. In the right display panel of Figure 1 (a), we show the IR result at a representative isovalue of the Argon Bubble data set.

DVR. For DVR, users specify the *transfer function* that maps scalar values to colors and opacities for direct rendering without using intermediate surface geometry. We implement the GPU-based raycasting [13] for interactive DVR. To cast a ray, we need to know where it enters and exits the volume, along which direction it travels, and how many samples we should take. We precompute the ray-plane intersections and store them into a texture as a prepass for rendering. Specifically, we first render the front face of the volume into a front-face texture, which is a screen-size texture encoding for every pixel, the entry point’s texture coordinate of the ray. In this prepass rendering, we also render the depth of entry points into a depth texture to maintain the correct depth order when the cutting plane is in use. To render the volume, we draw a full-screen quad with a raycasting shader. The shader looks up the front-face texture to find the pixel that need to be cast through, the texture coordinate, and the depth of the entry point. The ray direction is given by the vector from the entry and exit points in the texture space. Each sample’s position along the ray direction is computed via linear interpolation. In terms of how many samples we should take, we set the step size as half of a voxel. Users can adjust the **Sampling Rate** parameter to change the number of samples taken along each ray. Note that the assigned opacity also depends on the sampling rate. For example, when using fewer slices, the

opacity has to be scaled up, so that the overall intensity of the rendering results remains the same. We use Equation 1 to correct the opacity whenever users change the sampling rate s from the reference sampling rate s_0

$$A = 1 - (1 - A_0)^{s_0/s}, \quad (1)$$

where A_0 is the *original* alpha value, and A is the *corrected* alpha value. At each step along the ray, we sample scalar values from the texture and blend them front to back according to Equations 2 and 3

$$C_i = (1 - A_{i-1}) \times C_i + C_{i-1}, \quad (2)$$

$$A_i = (1 - A_{i-1}) \times A_i + A_{i-1}, \quad (3)$$

where C_i and C_{i-1} are, respectively, the colors at samples i and $i - 1$, and A_i and A_{i-1} are, respectively, the opacities at samples i and $i - 1$. With front-to-back compositing, we can terminate raycasting early if the color saturates (i.e., $A_i \geq 1.0$). The two rendering passes both require indexing into a 3D texture, which is supported by WebGL 2.0. In the left display panel of Figure 1 (a), we show the DVR result of the Argon Bubble data set.

4.2 User Interaction

VolumeVisual provides two side-by-side panels to show DVR and IR results along with a set of controls for users to explore VolVis results of different data sets. Major interactions provided in the interface fall into four categories: transfer function, cutting plane, lighting parameters, and viewing parameters.

Transfer function. DVR uses the transfer function to highlight features in the volumetric data by mapping scalar values to optical properties (colors and opacities). The user interface of the transfer function is an indispensable part of interactive design for DVR. Usually, the interface consists of a 1D curve editor for specifying the transfer function via a set of control points. VolumeVisual implements the most straightforward and commonly used 1D transfer function scheme. Color and opacity are assigned by the transfer function separately. For correct rendering, color needs to be multiplied by the corresponding opacity. VolumeVisual’s transfer function interface consists of three components, the color and opacity map, a set of control points, and a color picker. The color and opacity map is rendered by WebGL. Control points on the map are rendered by D3.js. The color picker is also implemented using D3.js, where users can choose a color for the corresponding control point specified in the color and opacity map. Figure 1 (a) shows the DVR display panel (left) along with the transfer function editor (Figure 1 (a)-1) and the color picker (Figure 1 (a)-2).




Figure 2: Different transfer functions emphasizing different regions of interest in the Supernova data set.

Note that both DVR and IR share the same color transfer function editor. For IR, the normalized isovalue is used as a texture coordinate to index into the color transfer function and assign a color for the isosurface. This helps users identify the given isovalue in DVR. VolumeVisual provides users with a default transfer function for each data set. Users can also adjust the transfer function in their preferred way; mapping different values ranges to different colors and opacities. Figure 2 shows DVR of the Supernova data set with different transfer functions.

Cutting plane. VolumeVisual enables the axis-aligned cutting planes for users to clip the volume along each spatial dimension. For each dimension, two sliders allow users to cut the volume in two opposite directions (Figure 1 (a)-3).

Both DVR and IR panels are synchronized. For example, in Figure 3 (a) and (b), the volume is clipped by the $Y(\max)$ and $Y(\min)$ sliders along the y -axis. By cutting through the volume, users can better observe the inside structure of the volume and focus on certain regions of interest. In the quiz component, students are required to use, in conjunction with the transfer function editor, the cutting planes to reveal a slab of the volume or enclose certain features or areas of the data set. This helps strengthen students' understanding of the basic VolVis concepts so that they can learn how to utilize DVR and IR to explore the volumetric data more efficiently.




Figure 3: Cutting along the y -axis through the Northridge Earthquake data set in (a) DVR and (b) IR. DVR results of the Bonsai data set (c) with and (d) without illumination.

Lighting parameters. Illumination is essential in DVR for improving the quality of rendering and creating realistic volumetric effects. VolumeVisual modifies the Blinn-Phong model to add important visual cues to the rendering. Per-pixel normal information is required to compute the local illumination with the Blinn-Phong model. In our implementation, we use central differencing to obtain the gradient vector at each voxel. For correct rendering, the ray direction needs to be transformed into the texture space. Desirable illumination provides not only high-quality rendering results but also extra information, such as gradient magnitudes. Figure 3 (c) and (d) show how illumination could affect the DVR results of the Bonsai data set. With the Blinn-Phong model, only those regions with high gradient magnitudes are shaded. As shown in Figure 1 (b)-4, the Lighting Parameters group provides sliders that enable users to tweak Ambient, Diffuse, Specular, Shininess, and Contrast sliders. The first four are the regular terms in the lighting model, while Contrast is the wrap lighting term that approximates scattering. Wrap lighting modifies the diffuse function so that the lighting wraps around the object beyond the point where it would normally become dark. This reduces the contrast of the diffuse lighting, which decreases the amount of ambient and diffuse lighting that is required.

Viewing parameters. VolumeVisual allows users to translate, rotate, and scale via mouse, touchpad, and keyboard. Alternatively, we also provide sliders for users to translate and scale the volume. Users can use their preferred ways to interact with the volume and gain a comprehensive understanding from different viewing directions. Both IR and DVR results are generated from the same volumetric data set and under the same view. This would help users better make side-by-side comparison and grasp the pros and cons of each method. Figure 1 (b)-5 shows Viewing Parameters unfolded in the interface. Two sliders allow translation along the x and y -axes, and one slider allows scaling the volume. Note that VolumeVisual does not provide translation along the z -axis because the Scale parameter could produce the same effect by zooming in the volume.

Other interactions. As shown in Figure 1 (c)-6, VolumeVisual provides additional controls. The Sampling Rate and Isovalue controls are specifically provided for DVR and IR, respectively. By adjusting the sampling rate, users can change the number of samples taken along the ray direction in volume raycasting, which could affect the resulting rendering quality. Figure 4 (a) and (b) show the rendering results of the Vortex data set with two different sampling rates. The Isovalue control allows users to switch between different representative isosurfaces of the same data set. Examples are shown in Figure 4 (c) and (d).

4.3 Quiz Component

VolumeVisual consists of two components: a study component and a quiz component. Figure 1 shows the interface of the study component of VolumeVisual. The study component is for students to learn general knowledge about




Figure 4: DVR of the Vortex data set with different sampling rates (a) 1.0 and (b) 5.0. IR of the UNC Brain data set with different isovalues (c) 22.4 and (d) 98.9.

VolVis techniques and get familiar with the VolumeVisual interface. In the quiz component, while a few questions test student's general knowledge about VolVis techniques, most of the questions check whether or not students can use the provided interactions to explore new data sets that do not appear in the study component. A question may be a textual one, including a text box for answering, or a single- or multiple-choice one, which is associated with a corresponding display panel. For data set related questions, the interface is quite similar to that of the study component. The difference is that now the question replaces the right display panel for showing IR. Some questions require viewing either the DVR or IR result in the single display panel. We add a switch to enable that, as shown in Figure 5. In the control panel, we remove the menu for data set selection, as only a single data set is associated with a question. The interface between each data set related question is slightly different depending on the question being asked. For such a question, we enable all the interactions in the Control Panel because this would test students to figure out the appropriate controls to get the answer. For the example shown in Figure 5, students are asked to adjust the lighting parameters in the Control Panel to generate a rendering result similar to the given reference image.




Figure 5: The screenshot of VolumeVisual (quiz component). Compared with the study component shown in Figure 1, we mainly replace the right IR panel with a panel for Q&A.

5 User Study

We conducted a user study to evaluate the effectiveness of VolumeVisual with two groups of participants, based on their familiarity with SciVis and VolVis concepts and techniques. Participants in the first group (VolVis-familiar

group) either have taken the visualization course before or are experts in SciVis and VolVis. In contrast, participants in the second group (VolVis-nonfamiliar group) have little prior knowledge about SciVis and VolVis. There are ten participants in the first group (including master students, Ph.D. students, and professors) and eight participants in the second group (including master and Ph.D. students). All the participants are from the respective computer science and engineering departments of two universities: University of Notre Dame and Clemson University.

The user study consists of two parts: an online survey and an online quiz. After the participants used the study component of VolumeVisual, they were required to complete a survey, which provides feedback about whether the study component improves their learning of VolVis techniques. The quiz questions testing the participants' general knowledge about VolVis and their ability to efficiently use all the interactions to explore the new data sets that are not seen in the study component. The quiz tool automatically records participants' answers, including all the parameters specified (e.g., lighting parameters, control points of the transfer function, etc.) for the data set related questions. To manually grade those questions that require participants to tweak the parameters, we developed a grading tool that loads the parameter setting for each question and renders the corresponding DVR and IR results.

5.1 General Information

The user study was conducted in three steps, and each step was five days apart. In the first step, each participant was sent introductory materials, including a tutorial video of VolumeVisual and slides introducing the background VolVis knowledge required for VolumeVisual. They also signed a consent form. To ensure that participants pay due attention to the subsequent survey and quiz, we paid each participant \$30 for completing the study. Participants had five days to go through the introductory materials before using the study component of VolumeVisual and completing the survey. In the second step, each participant was sent the URL of the VolumeVisual study component. A document explaining VolumeVisual was also provided. Participants practiced using the tool to familiarize themselves with VolumeVisual and finished the survey before moving to the quiz component. In the third step, each participant was required to take an online quiz. We then gathered the quiz results and performed the analysis.

5.2 Survey and Quiz Questions

The survey and quiz questions can be found in the Appendix. The survey consists of 20 rating questions with a five-point Likert scale ranging from "Strongly disagree" (1) to "Strongly agree" (5), and an open text question for general feedback. These rating questions are split into three categories: *user Interface* (S1-S14), *data sets* (S15-S18), and *general* (S19-S21). Out of the 21 quiz questions, four are single-choice, four are multiple-choice, six are open text questions, and the remaining seven are questions of other types. For questions of other types, participants were either required to use functions provided in Figure 1 to generate similar DVR results as the given image or to perform certain interactions such as *Cutting Plane* or *Lighting Parameters*. Among the 21 quiz questions, 12 are data set related, and the remaining nine are data set unrelated. Images accompany eight questions as a reference for participants to generate similar DVR results or answer questions. Data set unrelated questions test participants' general understanding of basic VolVis concepts. Data set related questions aim to examine whether participants can use the interface to explore the corresponding data set efficiently.

5.3 General Findings

Overall, we can report similar findings from both the VolVis-familiar group and the VolVis-nonfamiliar group. In either group, there was no significant quiz performance difference between master and Ph.D. students. Students (excluding the two professors) who had taken the visualization course before performed better than students who learned VolVis techniques on the fly in this study. Besides the monetary compensation, taking the quiz after using the study component could motivate participants to go through the introductory materials and study the concepts carefully. The survey results show that both groups were satisfied with VolumeVisual, and the overall feedback was positive. In the following, we describe the findings of the survey and quiz results in detail.

5.4 Survey Results

VolumeVisual was surveyed to obtain feedback from participants and identify ways for improvement. Table 2 shows the average scores (higher is better) and the standard deviations for the three types of questions. The results suggest that

	Data Set	User Interface	General	All Categories
Avg Rating	4.36	4.40	4.91	4.56
Std Dev	0.49	0.35	0.19	0.34
Avg Rating	4.22	4.35	4.31	4.29
Std Dev	0.38	0.34	0.56	0.43

Table 2: The average rating and standard deviation per survey question category and all categories. Top two rows: VolVis-familiar group. Bottom two rows: VolVis-nonfamiliar group.

VolumeVisual was positively received among both groups with average scores of 4.56/5.0 and 4.29/5.0, respectively. The VolVis-familiar group rated higher than the VolVis-nonfamiliar group. Upon closer inspection, this was largely due to the survey question S19: “*VolumeVisual provided an enjoyable way to learn about volume visualization.*” The average rating of this question from the VolVis-familiar group is 4.9, while that from the VolVis-nonfamiliar group is 4.43. This is likely because participants in the VolVis-familiar group have taken the visualization course and have already used other VolVis tools (such as VTK ParaView). They are familiar with VolVis, especially DVR, interactions, including editing the transfer function. So they could get started with VolumeVisual sooner than those in the VolVis-nonfamiliar group who just started using a VolVis tool from scratch. Although VolumeVisual has not yet been used in the classroom setting, the survey results suggest that it could help participants learn VolVis techniques by making the learning process more enjoyable.

One professor commented that an important aspect missing in VolumeVisual is a clear representation of the visualization pipeline, which could be seen in tools like ParaView or Inviwo [4]. He further commented that the information in VolumeVisual is all “at the same level”, which is not the case in reality. For example, when visualizing an isosurface, we start from the input data set, apply the marching cube algorithm [15], and then use the computer graphics pipeline for rendering. This kind of step-by-step organization is not easy to grasp in VolumeVisual. In response, we point out that we intend to hide the unnecessary complexity of VolVis techniques from participants, as the primary motivation for designing VolumeVisual is to make it suitable for beginners. VolumeVisual is one of the tools we create for teaching and learning the visualization course. Students will learn the entire VolVis concepts (including the pipeline of visualization) in class, and then VolumeVisual will be introduced to reinforce their understanding.

5.5 Quiz Results

With the quiz component, we aim to test participants’ general understanding of basic VolVis concepts and examine whether they can efficiently use the interface to explore the data sets. While grading choice- and text-questions is relatively easy, the rubrics for grading parameter-related questions are hard to control. This kind of question requires participants to tweak parameters (such as transfer function, viewing, and lighting parameters, etc.) to get good rendering or visualization results for the given data set. The difficulties come from the subjective nature of these questions because there are no standard criteria to measure whether or not a visualization for a data set is good. Difficult as it is, we still try to follow some rubrics while grading all the parameter-related questions. For example, for Q18 (refer to the Appendix), we first check if participants use the cutting plane to correctly enclose the hurricane’s eye and choose the color similar to the given reference image. These two steps are explicitly required in the question, and each step accounts for 30 points out of a total of 100 points (which is 60 points in total). For the remaining 40 points, we check if users adjust the lighting and viewing parameters ideally, for example, if the rendering could reveal gradient information about the data and if the volume is kept within an appropriate distance range (i.e., neither too far away nor too close from the camera). Although these two criteria are not explicitly required in the question, they are common sense for generating meaningful rendering results in computer graphics and visualization. Different parameter-related questions were graded with different rubrics, and we tried to maintain the same rubrics for grading the same question answered by different participants.

Table 3 and Table 4 show the average score and standard deviation for each of the 21 questions split by the two groups. The results suggest that both groups scored reasonably well with an average score of 88.98/100 and 76.50/100 for the VolVis-familiar group and VolVis-nonfamiliar group, respectively. The VolVis-familiar group outperformed the VolVis-nonfamiliar group, which is what we had expected. Although we had sent every participant the introductory materials and every participant had enough time to prepare for the upcoming quiz. The results suggest that taking the course along with VolumeVisual could still be a more effective way to learn VolVis than self-learning, as the stake (i.e., course grade) is higher for students taking the course. We can see the lowest scores in the VolVis-familiar group

Question	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Q11
VolVis-familiar group											
Avg Score	98.00	70.00	70.00	90.00	100.00	92.00	100.00	96.00	75.00	93.00	96.00
Std Dev	6.00	45.83	24.50	20.00	0.00	9.80	0.00	6.63	25.00	7.81	12.00
VolVis-nonfamiliar group											
Avg Score	97.50	50.00	68.75	75.00	93.75	67.50	75.00	95.00	50.00	86.25	73.75
Std Dev	6.61	50.00	34.80	35.36	16.54	28.17	43.30	7.07	0.00	4.84	33.52

Table 3: The average score and standard deviation per quiz question for Q1 to Q11.

Question	Q12	Q13	Q14	Q15	Q16	Q17	Q18	Q19	Q20	Q21	Q1-Q21
VolVis-familiar group											
Avg Score	94.00	100.00	95.00	100.00	59.60	89.00	84.00	98.00	81.00	88.00	88.98
Std Dev	8.00	0.00	8.06	0.00	25.02	15.78	11.14	6.00	12.21	29.93	13.03
VolVis-nonfamiliar group											
Avg Score	81.25	62.50	87.50	95.00	57.75	90.00	87.50	62.50	65.00	85.00	76.50
Std Dev	15.36	48.41	16.39	8.66	14.29	14.14	6.61	48.41	25.98	32.79	23.39

Table 4: The average score and standard deviation per quiz question for Q12 to Q21, and for all quiz questions (Q1-Q21).

come from Q2, Q9, and Q16, which examine participants’ knowledge about DVR techniques. For example, Q2 and Q9 test if participants are familiar with the raycasting process for DVR. In Q9, many participants confound the lighting parameters with the accumulated opacity in DVR and selected “ambient and diffuse values” as factors that may affect the final opacity. Similarly, in Q16, most participants overlooked the fact that DVR is designed for visualizing 3D scalar field data sets. Over half of the participants selected the “critical point” (which is a point in a vector field where the velocity is zero) that could not be shown in the DVR result of the tornado data set. Unsurprisingly, the scores of Q16 in the VolVis-familiar group are slightly higher. The standard deviation indicates that the performance of Q2 deviates more significantly than Q9 as Q2 is a single-choice question (i.e., the score is either 0 or 100). This is true for Q7, Q13, and Q19. Q7 and Q13 test fundamental knowledge about VolVis techniques. For example, Q7 asks, “Which process is not included in direct volume rendering?” The correct answer is “isosurface extraction”, which takes place before isosurface rendering. While all the participants in the VolVis-familiar group gave the correct answer to Q2, two participants in the VolVis-nonfamiliar group got zero, which significantly lifts the standard deviation. The clear discrepancy in these two groups’ scores is mainly due to the theoretical nature of these questions that are not directly related to VolumeVisual.

	Data	Non-Data	Text	Choice		Data	Non-Data	Text	Choice
Avg Score	90.05	88.89	93.50	84.07	Avg Score	79.71	72.22	86.04	65.03
Std Dev	10.24	12.66	4.68	15.94	Std Dev	11.42	17.05	9.14	13.49

Table 5: The average score and standard deviation for quiz per question type. Left: VolVis-familiar group. Right: VolVis-nonfamiliar group.

Table 5 shows the quiz results of VolVis-familiar group (left) and VolVis-nonfamiliar group (right) split by question type. One split evaluates the difference between data set related and data set unrelated questions. There is a noticeable gap in score between the two groups due to the difference in familiarity with VolVis techniques. A similar pattern in the two groups is that participants scored slightly higher in data set related questions than in data set unrelated questions. Note that even in the data set related questions that do not require participants to tweak the rendering parameters, we still enable all the interactions (except the transfer function editor) to allow them to explore the data set. This result suggests that when participants become familiar with data exploration using VolumeVisual, enabling all the interactions in data set related questions may result in a better quiz performance. However, this may not necessarily help them learn the theoretical background better. As the results suggest, participants in the VolVis-familiar group scored nearly 20 points higher in the choice questions than participants in the VolVis-nonfamiliar group. Upon closer inspection, two participants in the VolVis-nonfamiliar group only scored 33/100 in Q16 (a single-choice question testing participants’ general knowledge about DVR), which significantly lowers the average score. For the data set related questions, the score difference between the two groups is also apparent. For example, Q20 is the hardest data set related question because it requires participants to tweak all the parameters to generate a similar rendering as the

reference image. Six out of eight participants in the VolVis-nonfamiliar group scored lower than 80/100, while five out of 11 participants in the VolVis-familiar group got higher scores than 90/100. In general, the quiz results suggest that VolumeVisual does require some prior knowledge of VolVis to be used successfully by self-learners.

6 Conclusions and Future Work

We have presented VolumeVisual, a web-based educational tool to assist in teaching and learning VolVis concepts and techniques. VolumeVisual supports DVR and IR of various data sets, and provides two components (study and quiz). The effectiveness of VolumeVisual is demonstrated by conducting a formal user study with two participant groups (VolVis-familiar group and VolVis-nonfamiliar group). In the future, we would like to improve VolumeVisual by allowing the rendering of DVR and IR in the same view. This will help students better compare each technique's pros and cons in a superimposed view. According to the feedback to Q21 of the survey, another direction for improvement is to allow users to upload their own data sets. This feature will potentially enable VolumeVisual to engage more students with diverse backgrounds.

VolumeVisual is part of our effort to develop a series of educational tools for teaching and learning data visualization. We have released VolumeVisual online <https://sites.nd.edu/chaoli-wang/visvisual>, and the source code can be downloaded from <https://github.com/4321baoxueyi/volume-visual>. Besides previously developed FlowVisual [23, 22] and GraphVisual [10], we will further develop other tools for scientific visualization and information visualization, under the umbrella of VisVisual. Each tool of VisVisual can be used independently, either for classroom teaching and learning or self-learning.

Acknowledgements

This research was supported in part by the U.S. National Science Foundation through grant DUE-1833129.

References

- [1] <https://www.grc.org/visualization-in-science-and-education-conference/2015/>.
- [2] <https://vtk.org/>.
- [3] <https://www.paraview.org/>.
- [4] <https://inviwo.org/>.
- [5] M. Bostock, V. Ogievetsky, and J. Heer. D³ data-driven documents. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2301–2309, 2011.
- [6] S. Bruckner and T. Möller. Isosurface similarity maps. *Computer Graphics Forum*, 29(3):773–782, 2010.
- [7] C. Chen. Top 10 unsolved information visualization problems. *IEEE Computer Graphics and Applications*, 25(4):12–16, 2005.
- [8] J. Fekete. The InfoVis toolkit. In *IEEE Symposium on Information Visualization*, pages 167–174, 2004.
- [9] J. Hertzberg and A. Sweetman. A course in flow visualization: The art and physics of fluid flow. In *Proceedings of American Society for Engineering Education Annual Conference*, 2004.
- [10] M. Imre, W. Chang, S. Wang, C. Trinter, and C. Wang. GraphVisual: Design and evaluation of a web-based visualization tool for teaching and learning graph visualization. In *Proceedings of American Society for Engineering Education Annual Conference*, pages 28.501.1–28.501.15, 2020.
- [11] M. Imre, J. Tao, and C. Wang. Efficient GPU-accelerated computation of isosurface similarity maps. In *Proceedings of IEEE Pacific Visualization Symposium (Visualization Notes)*, pages 180–184, 2017.
- [12] J. Kniss, S. Premoze, C. D. Hansen, and D. S. Ebert. Interactive translucent volume rendering and procedural modeling. In *Proceedings of IEEE Visualization Conference*, pages 109–116, 2002.
- [13] J. H. Krüger and R. Westermann. Acceleration techniques for GPU-based volume rendering. In *Proceedings of IEEE Visualization Conference*, pages 287–292, 2003.

- [14] W. Li, K. Mueller, and A. E. Kaufman. Empty space skipping and occlusion clipping for texture-based volume rendering. In *Proceedings of IEEE Visualization Conference*, pages 317–324, 2003.
- [15] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. In *Proceedings of ACM SIGGRAPH Conference*, pages 163–169, 1987.
- [16] F. Paas, A. Renkl, and J. Sweller. Cognitive load theory and instructional design: Recent developments. *Educational Psychologist*, 38(1):1–4, 2003.
- [17] S. Papert and I. Harel. Situating constructionism. *Constructionism*, 36(2):1–11, 1991.
- [18] D. N. Perkins. Technology meets constructivism: Do they make a marriage? *Educational Technology*, 31(5):18–23, 1991.
- [19] M. Stieff, R. C. Bateman, and D. H. Uttal. Teaching and learning with three-dimensional representations. In J. K. Gilbert, editor, *Visualization in Science Education*, pages 93–120. Springer, 2007.
- [20] Y. Tanahashi, N. Leaf, and K.-L. Ma. A study on designing effective introductory materials for information visualization. *Computer Graphics Forum*, 35(7):117–126, 2016.
- [21] K. L. Vavra, V. Janjic-Watrich, K. Loerke, L. M. Phillips, S. P. Norris, and J. Macnab. Visualization in science education. *Alberta Science Education Journal*, 41(1):22–30, 2011.
- [22] M. Wang, J. Tao, J. Ma, Y. Shen, and C. Wang. FlowVisual: A visualization app for teaching and understanding 3D flow field concepts. In *Proceedings of IS&T Conference on Visualization and Data Analysis*, pages 476–1–476–10, 2016.
- [23] M. Wang, J. Tao, C. Wang, C.-K. Shene, and S. H. Kim. FlowVisual: Design and evaluation of a visualization tool for teaching 2D flow field concepts. In *Proceedings of American Society for Engineering Education Annual Conference*, pages 23.609.1–23.609.20, 2013.
- [24] R. Westermann and T. Ertl. Efficiently using graphics hardware in volume rendering applications. In *Proceedings of ACM SIGGRAPH Conference*, pages 169–177, 1998.

A Appendix

A.1 Survey Questions

User Interface

- S1 The two side-by-side panels of direct volume rendering and isosurface rendering help me better understand the pros and cons of these two different volume visualization techniques.
- S2 The layout of the VolumeVisual interface is appropriate.
- S3 The views of direct volume rendering and isosurface rendering could be easily manipulated by mouse (touchpad) and keyboard.
- S4 The cutting plane panel is easy to use and helps me better observe the inside structure of the volume.
- S5 The lighting parameters panel is sufficient for getting desirable rendering quality of direct volume rendering and isosurface rendering.
- S6 The onboarding feature helps me better get oriented with VolumeVisual.
- S7 The “about” page of VolumeVisual provides enough background knowledge about VolumeVisual.
- S8 The overall interface design of VolumeVisual is intuitive and easy to follow.
- S9 For each data set, the default settings (transfer function, lighting and viewing parameters, etc.) for all the parameters are ideal.
- S10 The color picker provides enough color options to generate preferred transfer functions.
- S11 The visual effects for user interface controls help me better distinguish different controls.
- S12 The interactions are sufficient for understanding the basic volume visualization concepts.
- S13 The color picker and transfer function editor are easy to use for specifying color and opacity.
- S14 The interactions, along with the given hints, are easy to understand and follow.

Data Sets

- S15 The data sets provided are interesting and representative enough to help me better learn and differentiate direct volume rendering and isosurface rendering techniques.
- S16 The data sets provided are sufficient for me to learn volume visualization techniques.
- S17 The description of each data set is informative and clear.
- S18 The isovalues provided for all the data sets are sufficient and representative.

General

- S19 VolumeVisual provided an enjoyable way to learn about volume visualization.
- S20 The user interface of VolumeVisual is friendly and simple.
- S21 What one change to the items above, or other parts of VolumeVisual, would improve your learning of volume visualization?

A.2 Quiz Questions

Text questions are labeled with a †, single-choice questions with a ★, multiple-choice questions with a ∼, and questions of other types with a ψ.

- Q1 † Comparing direct volume rendering and isosurface rendering, point out their respective pros and cons.
- Q2 ★ Assume the resolution of volume is 128 in each dimension. The sampling rate is one voxel wide during the raycasting process. How many samples at most are taken along one ray?
- Q3 ∼ Which operations will likely increase the computation cost of direct volume rendering?
- Q4 † Turn on the axis (red, green, and blue represent X , Y , and Z axis). Assuming X and Y represent East and North in the data set. Which two wind directions does this data set illustrate?

- Q5 ~ Which techniques are the most often used for optimizing the efficiency of direct volume rendering?
- Q6 ψ The basin of the bonsai data set is already shown in the rendering. Use the transfer function editor to show the leaves and branches. (You must click “Save” before moving to the next question.)
- Q7 $\star$ Which process is not included in direct volume rendering?
- Q8 ψ The rendering of the brain is upside down and too small. Adjust the viewing parameters to correct these. Use the color picker to choose natural colors for skin and brain.
- Q9 ~ Which options will influence the final accumulated opacity of direct volume rendering?
- Q10 ψ Adjust the lighting parameters and sampling rate to generate a rendering similar to the image below.
- Q11 $\dagger$ Given the two direct volume rendering images of the lobster data set below, why the right image has a better visual quality in rendering? Explain briefly.
- Q12 ψ Adjust the transfer function editor to generate a direct volume rendering result similar to the corresponding isosurface rendering result.
- Q13 $\star$ Which technique is most often used for isosurface extraction?
- Q14 ψ Adjust the lighting parameters and use the cutting planes to show a slab of the hurricane similar to the image below.
- Q15 $\dagger$ Given the two direct volume rendering images below, which one has a better transfer function applied? Please state at least two reasons.
- Q16 ~ Which information could not be read from direct volume rendering of the tornado data set?
- Q17 $\dagger$ Comparing the raycasting method and the texture-based slicing method for direct volume rendering, point out their respective pros and cons.
- Q18 ψ Given the isosurface rendering image below, use the cutting planes to reveal the hurricane’s eye in direct volume rendering. Also, choose a similar color for the eye.
- Q19 $\star$ In direct volume rendering, scalar values at intermediate locations are obtained by interpolating scalar values at neighboring voxels. How many 3D data texture lookups does this process perform in the trilinear interpolation of one intermediate location?
- Q20 ψ Tweak all the parameters provided in the interface to generate the direct volume rendering result similar to the image below.
- Q21 $\dagger$ The two images below are rendered by direct volume rendering and isosurface rendering, both of which highlight the same scalar value. Please judge which one is rendered by direct volume rendering and which one is rendered by isosurface rendering. Explain briefly.

Data sets used in quiz questions

- Q1 None
- Q2 None
- Q3 None
- Q4 Atmospheric
- Q5 None
- Q6 Bonsai
- Q7 None
- Q8 UNC Brain
- Q9 None
- Q10 Five Jets
- Q11 Lobster
- Q12 Combustion (OH)
- Q13 None
- Q14 Hurricane Isabel (VAPOR)
- Q15 Foot
- Q16 Tornado
- Q17 None
- Q18 Hurricane Isabel (WSMAG)
- Q19 None
- Q20 Combustion (HR)
- Q21 UNC Brain