

Context-Preserving Volumetric Data Set Exploration Using a 3D Painting Metaphor

L. Faynshteyn and T. McInerney

Dept. of Computer Science, Ryerson University, Toronto, ON, Canada, M5B 2K3

Abstract. This paper presents a technique for fast and intuitive visualization and exploration of unsegmented volumetric data sets. It combines a Maximum Intensity Difference Accumulation (MIDA) visualization algorithm with a 3D interaction model based on a painting metaphor. Under this model the Volume of Interest (VOI) is defined by “painting” an envelope enclosing features of interest. The key advantage of the model is the ability to interactively preview the effect of the painting operation before it is applied, allowing the user to both explore and paint in the volume space at the same time. The result is a flexible and intuitive technique that uses a single, consistent interaction style to achieve desired contextual views and perform several types of volume exploration and editing operations. Examples using several volumetric data sets and the results of a user study are presented to validate the effectiveness, flexibility and ease of use of the technique.

1 Introduction

Despite the significant amount of research in the field of volume rendering, when it comes to providing a user with easily understandable and predictable ways of specifying how the data is to be visualized, there is still much room for improvement. Traditionally this specification has been realized by means of a Transfer Function (TF). However, specifying a TF that achieves the desired visual effect still remains a non-trivial task, even for experienced users. This is primarily due to the fact that TF specification is usually performed in a histogram space and thus doesn’t provide the user with a natural feel for the correspondence of features in the histogram with features in the volumetric data set.

The challenge of finding a perfect TF is also compounded by the now common requirement of visualizing areas of interest in the volume using different sets of parameters than the surrounding region, in an effort to preserve the visual context. This secondary goal presents visualization challenges in its own right as navigation in the 3D space of a volume and specification of different VOIs in a 3D data set is often a complex operation. Among the more prominent issues are VOI depth specification and effective depth perception cues, as well as volumetric features separation and occlusion. The traditional approach to the VOI specification, feature separation, and occlusion problems is to use multiple 2D views along with a variety of 2D VOI editing capabilities. While certainly useful in many scenarios, when it comes to volume exploration and VOI specification,

multiple 2D views complicate the user interaction model and it is often difficult to understand the spatial relationships among all the views.

To meet these challenges, we propose a volume visualization and exploration model that combines the Maximum Intensity Difference Accumulation (MIDA) visualization method with an intuitive and flexible volume painting interaction model. The MIDA visualization technique was proposed by Bruckner et al. [1] as a means of quickly visualizing volumetric data sets without the necessity of specifying TFs. The MIDA technique uses a conventional Direct Volume Rendering (DVR) approach and modifies it to exploit inherent data characteristics of the underlying volumetric data set by modulating the accumulated opacity of the image in accordance with changes in the data intensity values. To control which features of a volumetric data set should be emphasized, only a range defined by the minimum and maximum data intensity values needs to be specified.

The painting interaction model supports free-form and constrained painting and editing of envelopes defining VOIs. Features of interest are surrounded by positioning the tip of a paint “brush” at the desired location and depositing a “blob” of paint, extending the envelope by merging with it. The envelope can be thought of as “thick paint”, implemented using Metaballs [2] (see section 3.2), and volumetric features inside and outside of the envelope can be visualized using two separate (but functionally identical) sets of MIDA visualization parameters. The key difference between our approach and many other techniques is that it allows the user to preview, in real time, the results of the painting operation before the new blob of paint is deposited. That is, the Metaball-based paint blob acts as a lens that can be moved around to visualize the volume interior before deciding on its final location. This separation of blob positioning and depositing helps the user to interactively guide the painting process and results in a high degree of control over the VOI shape definition.

The combination of the two techniques allows the user to operate within a single 3D data view and employs a minimum of visualization controls. It also enables a variety of possible volume data visualization, exploration and editing operations (e.g. contextual views, volume carving, “Magic Lantern”-like behavior [3] etc.) to be realized using a single, consistent interaction model. In addition, the proposed approach does not depend on any particular visualization method and other DVR techniques can be easily substituted should MIDA prove insufficient in terms of achievable visual results. A number of examples showcasing the visualization and exploration capabilities of the proposed approach are presented, along with the results of a user study validating the ease of use, intuitiveness and effectiveness of the approach.

2 Related Work

Over last two decades, a considerable amount of DVR research has focussed on improving TF specification. Since 1D TFs often suffer from poor volumetric feature separation, more sophisticated higher dimension TFs are now commonly used. Kniss et al. [4] propose a 2D TF specification process based on data in-

tensity and gradient magnitude values that enables better feature separation by representing material boundaries as arches on a 2D histogram. Correa and Kwan-Liu [5] use the relative size of volumetric features to construct a TF with improved feature classification. Sereda et al. [6] construct so called LH histograms that show lower and higher intensities of the materials inside a volume, resulting in a better definition of material boundaries. All these methods, however, force the user to either work in the feature space of 2D histograms or manipulate some other data driven parameters that control the TF specification process. This, in turn, usually results in a time consuming and tedious trial and error specification process.

To simplify the TF specification various automatic and semi-automatic techniques have been proposed. Kindlmann et al. [7] calculate the curvature of volumetric features using iso-surfaces, which is then used as a basis for designing a TF and to generate non-photorealistic images by highlighting the contours of volumetric structures. Prassni et al. [8] introduce a concept of Curve-skeletons, that are used to classify volumetric features inside a pre-segmented VOI using predefined shape descriptors. The resulting shapes are then merged in a process supervised by the user. Haidacher et al. [9] present a method of adaptively estimating statistical properties of the data intensity values in the neighborhood of each sampled point within a volume. These properties are used for TF specification which enables the differentiation of volumetric features. While these and other automatic approaches significantly reduce the complexities of the TF specification, they also require lengthy pre-processing steps.

To better facilitate understanding of the information contained in a volume, it is necessary to be able to explore structures hidden inside it. This can be achieved by specifying a VOI containing the structures of interest and then high-lighting these structures by using some sort of distinctive visual style. At the same time, parts of the volume surrounding the VOI should be preserved for contextual purposes. Most of the approaches that attempt to meet these goals are based on concepts borrowed from the field of traditional illustration. Among these are non-photorealistic rendering with emphasized silhouettes [7], stylized shading, contouring, transparency and haloing [10]. These methods help the user understand the data by making certain features within a volume more prominent. However, highlighting features of interest is often not sufficient since they can still be occluded by other structures.

Several context-preserving techniques addressing the problem of occlusion have been proposed. Inspired by the idea of ghosted views, Bruckner et al. [11] propose a method that uses a shading intensity function for opacity modulation of structures within a volume. Another approach by Bruckner et al. [12] use compositing and masking to combine several pre-rendered layers and draw the user's attention to certain parts of the image while preserving the overall context. The use of masking also enables the effect known as a Magic Lens or Lantern, which is another popular context preserving technique. Monclus et al. [3] use a second TF to visualize a cone-shaped VOI in a way that is distinct from the surrounding material. Various standard and custom magnification effects, within

a user defined arbitrarily-shaped 2D region on the screen, are employed in the work by Wang et al. [13]. Bruckner and Gröller [14] use a 3D volumetric painting approach, where the user can define a surface-bound 3D envelope that encloses the features of interest, which are then visualized using an alternative predefined TF to provide contextual data representation.

Another way of dealing with the problem of occlusion is based on the idea of physically removing the occluding parts of the volume or deforming parts of the volume in such a way as to make the features of interest visible. The work presented by Bruckner and Gröller [15] and Ruiz et al. [16] is based on the concept of exploded views. Weiskopf et al. [17] cut away parts of a volume by using various geometric primitives and depth test algorithms. Bernhard et al. [18] use a deformable cutting plane with an adjustable sphere of influence to define the VOI and exclude all other volumetric regions. User-controlled drilling, lasering and peeling operations are used by Chen et al. [19] to remove parts of the surface of a volumetric object to reveal the underlying structures.

Upon considering the complexities and limitations of the TF-based approaches, we decided to focus on algorithms and interaction models with very simple and easily extendible interfaces. The non-TF based MIDA approach combined with the flexible Metaball-based envelope painting forms the basis of our system. Our envelope painting approach shares similarities with many of the techniques mentioned above. In particular, the surface painting described in [14] is similar to our surface painting mode, and the carving and peeling operations presented in [19] can be used to achieve visual results similar to some of ours. However, unlike most of these techniques, we use a higher-level shape representation of the envelope to visualize and define the VOI (i.e. Metaballs). Instead of using sets of independent surface-bound volumetric brushes [14], arbitrarily-shaped depth-fixed peeling regions [19], or shape-restricted envelopes [3], we propose a more general approach that combines and extends all of these visualization and interaction capabilities in a single, flexible and consistent model. The key feature, real time previewing, allows for “magic lens”-like behavior in both the 2D and 3D painting modes and enables precise and easy envelope definition and editing capabilities. The ability to easily control the range of sampled intensity values in the MIDA technique, combined with an unrestricted 3D envelope definition mechanism, enables flexibility in terms of possible types of volume cutting and peeling operations. Furthermore, recognizing that the MIDA visualization technique might not provide the desired visual results in all cases, our system was also designed and implemented in such a way as to be independent of the particular DVR technique used, allowing for easy and straightforward integration with TF-based visualization models.

3 Methodology

3.1 Maximum Intensity Difference Accumulation

MIDA works by altering the monotonically growing opacity accumulation behavior, inherent to the traditional DVR approach, by modulating any previ-

ously accumulated opacity values by the amount of positive difference between the currently sampled data intensity value and the previously encountered maximum intensity value. This, in turn, allows volumetric structures represented by higher intensity values (e.g. bones vs. skin) to be given more prominence in the resulting image. A simple linear mapping is used to map data intensity values to gray-scale color and opacity values. The obtained gray-scale values can also be modulated by a user specified color value, resulting in more realistic visualizations. In addition, and rather importantly, a range of mapped data intensity values can be specified. This allows the user to easily control (or filter), with just one set of GUI slider controls, which volumetric features will contribute to the final image and which will not (Fig. 1). As will be shown later, we use this feature extensively both in our exploration work-flow and to provide volume carving and peeling operations.

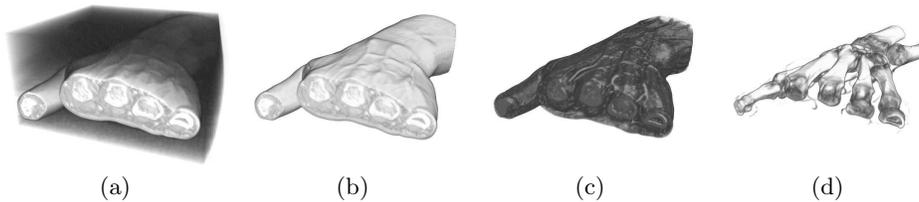


Fig. 1. By simply adjusting the range of mapped data intensity values and specifying the modulation color, the user is able to control the visualization process: (a) an unfiltered and uncolored volume visualization; (b-d) different volumetric features are visible and colored according to the user’s preference.

3.2 3D Envelope Painting Overview and Work-flow

We employ a single interaction model that can be used for volume exploration, VOI definition, and contextual data view generation. It is based on a painting metaphor and realizes a concept of “thick paint” that enables the user to select features of interest within a volume by painting an envelope either in a 2D space of the screen or a 3D space of a volume. The painting process is controlled by moving a virtual brush tip in the 3D space of a scene. The brush is represented by a blob of paint (centered at the brush tip position) that can be smoothly merged with the current envelope painted thus far, allowing for a continuous envelope extrusion. Since everything inside the paint blob (and current envelope) is instantly visualized using a separate MIDA visualization style, the user can move the blob around the current region and preview the volume interior. When the user is satisfied with the results of this previewing operation, they “deposit” the blob with a mouse click, extending the envelope (i.e. spreading the paint)(Fig. 2). This previewing capability can also be used as a tool for volume exploration, acting as a sort of “flash light” or lens, allowing the user to quickly look around the volume interior.

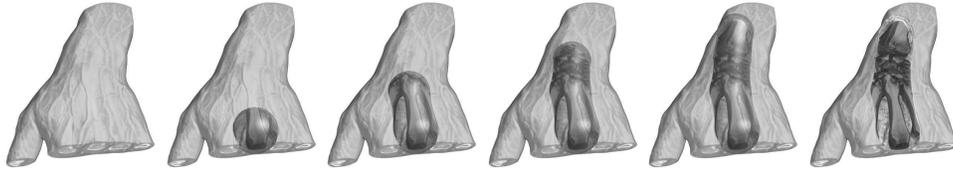


Fig. 2. Envelope painting depicted as a sequence of steps (from left to right) for previewing the results of the painting operation by moving the brush tip to new locations, and spreading the paint (i.e. painting the envelope) by depositing new blobs of paint at desired locations. When the desired features are painted, the envelope is made transparent (rightmost image).

3.3 Painting Modes and Volume Exploration Operation Types

One of the advantages of the Metaball-based painting technique is that the paint blobs can be moved around in the 3D space of the volume, in screen space, or can slide around on the surface of the objects in the volume. To paint in the 3D space using a 2D input device (i.e. a mouse or a touch screen) we allow the user to paint only in the space of a 2D painting plane that is always perpendicular to the current viewing direction of the 3D scene. By changing the position of the mouse (or finger etc.) in the screen space coordinates, the user defines the position of the brush tip as projected onto the painting plane. To add a third dimension to the painting process, the position, orientation and depth of the painting plane relative to the volumetric object can be changed at any time. This constrained interaction of adjusting the orientation and position of the painting plane relative to the volumetric object and of depositing blobs of paint to define arbitrarily-shaped 3D envelopes keeps the 2D input device based user interaction model simple, intuitive and consistent.

3D Surface Painting Mode: In this mode the user can explore a volumetric object by automatically following its surface with a brush tip (Fig. 3b). This is done by overriding the depth of the brush tip, and hence the depth of the painting plane at the current projected position, by the depth value of the object’s surface. In addition, an alternative brush type is available in this mode. This type of brush is based on a superquadric metaball (as opposed to a regular quadric metaball) to define cylinder-shaped brushes that are used to paint envelopes that more closely conform to the shape of a volumetric object (e.g. a “thin shell” envelope can be defined that surrounds, for example, only a skin layer) (Fig. 3c). The orientation of the cylinder-shaped brush is based on the surface normal orientation. The depth of the object surface penetration can be controlled by the user.

2D Screen Painting Mode: In this mode the depth of the painting plane is constant and any depth information related to the envelope is discarded. Instead, the outline of the shape of the envelope is used as a stencil, inside of which the secondary MIDA visualization style is applied. Since there is no depth specified, the effect is applied all the way through the volume in the direction of projected rays (Fig. 3d). In addition, unlike the 3D painting modes where the painted

envelope always maintains position and orientation relative to the volumetric object, in the 2D screen painting mode the painted envelope stays glued to the screen regardless of the object’s orientation in the scene. This allows the user to examine the object from different perspectives using the 2D envelope as a sort of arbitrarily-shaped “X-ray device” (or previously described Magic Lantern), that can highlight volumetric features while the object’s orientation is being changed.

Brush and Envelope Related Operations: In all painting modes, the size (radius) (Fig. 3e) and color of the brush can be adjusted as well as the color and transparency of the envelope. Finally, the shape of the envelope can be edited by using an “eraser” brush. This brush simply erases parts of the envelope that fall within its sphere of influence (Fig. 3f). If desired, the envelope can also be completely discarded.

Volume Editing: Another set of volume exploration operations is based on the previously mentioned ability to easily control the range of mapped data intensity values in the MIDA visualization technique. By setting the range to zero, the user can discard (carve or peel) parts of the volume that correspond either to those contained inside of the envelope or outside of it. That is, the paint brush becomes a sculpting tool providing a very simple, yet powerful and flexible method for volume exploration and context-preserving visualization (Fig. 3g,h).

Alternative rendering modes: In cases when MIDA proves to be insufficient in terms of desired visual results, this rendering mode can be easily substituted with alternative DVR methods, while preserving all of the original painting modes and functionality. To showcase this capability, we have implemented additional 1D and 2D TF-based rendering modes that can be used to visualize the parts of the volume both inside and outside of the envelope (Fig. 3i).

3.4 Implementation Details

In the current implementation, the brush and envelope are realized using the technique for defining implicit 3D surfaces known as Metaballs [2]. Metaballs are capable of merging with each other based on their proximity, radius and the chosen field function, creating organic-looking and flowing shapes. We use two types of metaballs: quadric, corresponding to a conventional spherical brush (i.e. a blob), and superquadric, used to define disk or cylinder shaped brushes. The six-degree Wyvill polynomial [20] is used to model the metaball field falloff function. To convert the implicitly defined envelope surface into a 3D polygonal representation, we use the Marching Cubes algorithm. A fast GLSL shader implementation of the algorithm is used to extract the iso-surface and then generate and capture the resulting geometry into a memory vertex buffer (which can be reused later, removing the need to re-compute the geometry for every frame). We use Front-to-Back Depth Peeling to combine polygon-based (i.e. the envelope) and volumetric geometry rendering. The method works by rendering the polygonal geometry in the scene layer by layer, as viewed from the camera position, and filling in the gaps between the layers with the results of the DVR. The results are blended together and there is no limit on the number of rendered layers. Furthermore, as was mentioned earlier, an additional advantage of this

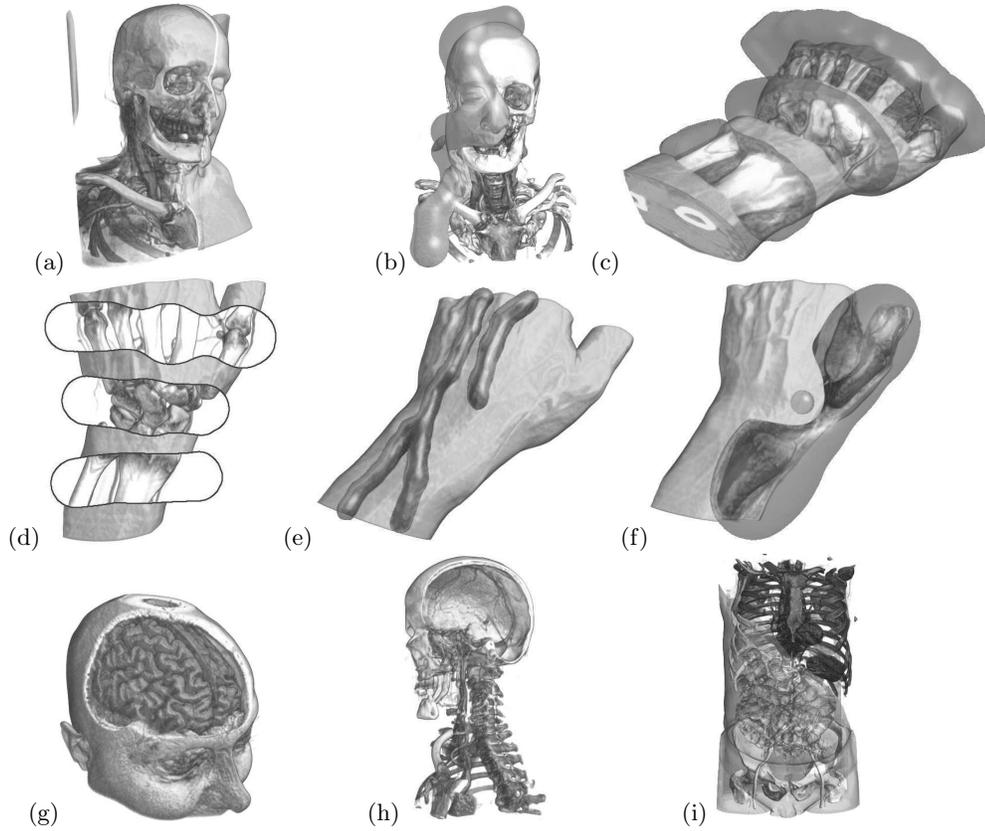


Fig. 3. (a) 3D free-hand painting is used to envelope features deep inside the volume and achieve the desired contextual view; (b,c) 3D surface painting, with spherical and surface aligned cylinder-shaped brushes, slide over the objects' surface; (d) an example of 2D screen painting mode; (e) a small sized spherical brush is used for precise painting on the skin surface; (f) an eraser brush is used to edit the envelope; (g,h) the peeling and carving operations are performed by setting the MIDA mapped data ranges to 0; (i) the MIDA and 2D TF-based DVR techniques are used together to visualize different parts of the same volumetric object.

approach is its independence of the type of DVR technique used, allowing for the easy extensibility of the visualization model. Gaussian-smoothed position and normal maps, corresponding to the surface of the visualized object, are generated on-the-fly and are used to control the depth and orientation of the brush in the 3D surface painting mode. For implementation details refer to [21].

4 Results and Validation

In order to evaluate the effectiveness of the proposed exploration technique, as well as to gauge how intuitive and easy to learn and control the painting-based interaction model is, a user study was conducted with ten male subjects and one female subject. Ten of the subjects were undergraduate students and one subject was a graduate student, all of whom were from either the department of computer science or engineering. All subjects were first-time users of the implemented exploration technique and none had any previous experience with any medical software. Approximately half, however, did have limited experience with various 3D modeling programs (i.e. Blender, Maya, etc.). In order to better gauge the level of the participants' proficiency performing navigational tasks in virtual 3D environments, they were also asked to state the number of Hours per Week (HPW) spent using a mouse and playing video games, with the average numbers coming to 25.5HPW and 11.5HPW respectively. The participants were

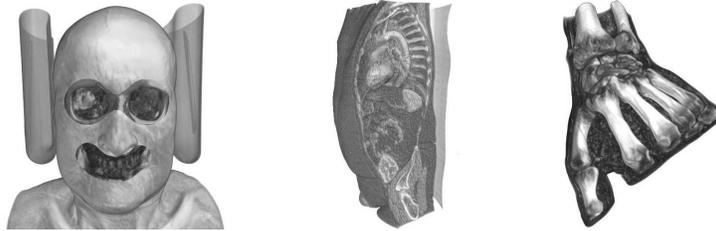


Fig. 4. Examples of pre-rendered reference images from each set of trials. Left: an example of the 2D painting task, middle: an example of the 3D carving/peeling task and, right: an example of the 3D context-preserving task. The subjects were asked to match these images.

asked to use the implemented software to perform several tasks. A total of nine trials, each based on a different data set, were performed. The goal of each trial was to achieve an approximate visual match to a pre-rendered reference image (4). The trials were further broken down into 3 sets: three 2D painting trials, three 3D volume carving trials and three 3D context-preserving trials. In the course of the trials from both the 3D carving/peeling and 3D context-preserving sets, the users had to use both the free-hand and the surface painting modes. Before each set of trials each participant was given instructions using a separate demonstration data set and was also given time to play with the software to

familiarize themselves with the visualization controls and the painting mechanism. The time for each trial was recorded for statistical purposes, although the participants were informed that achieving a visual match in the shortest possible time was not an objective. The average measured trial completion time for these non-expert users was 54 seconds, with results varying from 15 seconds in the case of simple 2D screen painting tasks and up to 2 minutes for more complex 3D painting tasks. Upon completion of all trials the participants were asked to fill out the questionnaire and indicate their level of agreement or disagreement with each statement using a 7-point Likert scale (5).

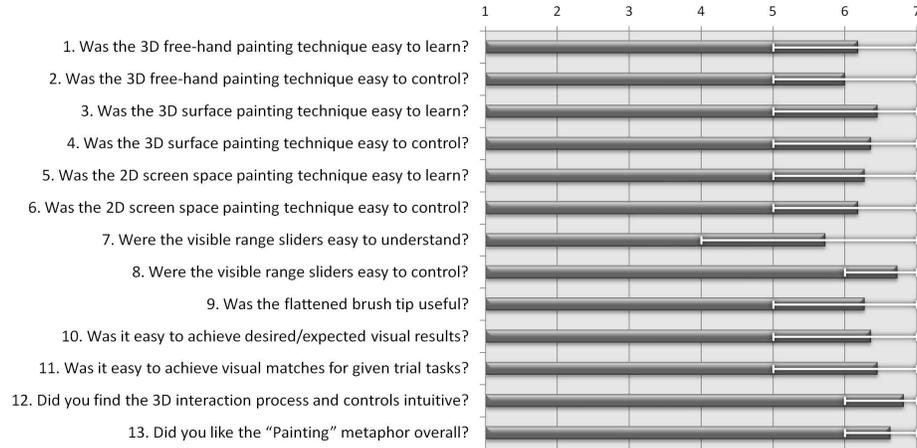


Fig. 5. The results of the user study questionnaire based on a 7-point Likert scale.

As is evident from Figure 5, the results were very positive, with a mean average of 6.0 or more. In addition, as can be seen from the spread of the marks, there was little variability in the results. The only question that received a mean below 6.0 was related to the sliders that control the visible range of values for the MIDA visualization mode. Another common comment from participants was that, even though gauging the depth of the brush tip in relation to the volumetric structures in the scene was usually not a problem, under certain conditions it still could be somewhat confusing, and thus they would prefer to have additional depth cues. We plan to address both of these shortcomings in the near future.

5 Conclusion

As evidenced by ongoing research, providing an intuitive way of visualizing volumetric data sets with predictable visual results is still a challenge. The vast majority of the current techniques rely on a TF mechanism to map volume data characteristics to the visual domain. However, TF specification often remains a

complex and time consuming trial and error process. The semi-automatic and automatic approaches address this problem with varying degrees of success and often at the expense of interactivity due to the costly pre-processing operations.

Another challenge lies in providing the user the ability to efficiently navigate the volumetric space and to easily select features of interest while preserving contextual data views. Various techniques based on the ideas of non-photorealistic rendering, a magic lens metaphor, painting, volume sculpting and peeling, and others have been proposed by researchers. To meet these two challenges, we have proposed a volume visualization and exploration framework that combines a non-TF based MIDA DVR method with a simple, yet flexible Metaball-based painting technique. This synthesis enables many of the previously described context-preserving volume exploration and editing techniques to be realized through a single, consistent and intuitive interaction model. The results of a user study as well as feedback the users confirm the validity of these claims and the viability of the proposed solution for relatively noise-free data sets. Future work includes developing more sophisticated undo/redo operations and envelope editing operations by taking advantage of the Metaball-based envelope shape representation. Allowing the user to paint additional VOIs with the ability to specify the visualization parameters for each VOI independently should greatly enhance the exploration capabilities of the system. Finally, and perhaps most importantly, we are currently working on the incorporation of several painting-based interactive segmentation algorithms into the work-flow in order to allow for the visualization and exploration of noisy data sets. Rather than attempting to construct complex TFs, the idea is to retain and augment our simple but flexible painting interaction model and use it to control and constrain the segmentation algorithms.

References

1. Bruckner, S., Gröller, M.: Instant volume visualization using maximum intensity difference accumulation. *Computer Graphics Forum* **28** (2009) 775–782
2. Blinn, J.F.: A generalization of algebraic surface drawing. *ACM Transactions on Graphics* **1** (1982) 235–256
3. Monclús, E., Díaz, J., Navazo, I., Vázquez, P.: The virtual magic lantern: an interaction metaphor for enhanced medical data inspection. In: *ACM Symposium on Virtual Reality Software and Technology*, Kyoto, Japan (2009) 119–122
4. Kniss, J., Kindlmann, G., Hansen, C.: Multidimensional transfer functions for interactive volume rendering. *IEEE Trans. on Visualization and Computer Graphics* **8** (2002) 270–285
5. Correa, C., Kwan-Liu, M.: Size-based transfer functions: A new volume exploration technique. *IEEE Trans. on Visualization and Comp. Graphics* **14** (2008) 1380–1387
6. Sereda, P., Bartoli, A., Serlie, I., Gerritsen, F.: Visualization of boundaries in volumetric data sets using lh histograms. *IEEE Trans. on Visualization and Computer Graphics* **12** (2006) 208–218
7. Kindlmann, G., Whitaker, R., Tasdizen, T., Moller, T.: Curvature-based transfer functions for direct volume rendering: methods and applications. In: *14th IEEE Conference on Visualization 2003 (VIS'03)*, Seattle, WA (2003) 513–520

8. Prassni, J., Ropinski, T., Mensmann, J., Hinrichs, K.: Shape-based transfer functions for volume visualization. In: IEEE Pacific Visualization Symposium 2010 (PacificVis '10), Taipei, Taiwan (2010) 9–16
9. Haidacher, M., Patel, D., Bruckner, S., Kanitsar, A., Gröller, M.: Volume visualization based on statistical transfer-function spaces. In: IEEE Pacific Visualization Symposium 2010 (PacificVis '10), Taipei, Taiwan (2010) 17–24
10. Bruckner, S., Gröller, M.: Style transfer functions for illustrative volume rendering. *Computer Graphics Forum* **26** (2007) 715–724
11. Bruckner, S., Grimm, S., Kanitsar, A., Gröller, M.: Illustrative context-preserving exploration of volume data. *IEEE Transactions on Visualization and Computer Graphics* **12** (2006) 1559–1569
12. Bruckner, S., Rautek, P., Viola, I., Roberts, M., Sousa, M., Gröller, M.: Hybrid visibility compositing and masking for illustrative rendering. *Computers & Graphics* **34** (2010) 361–369
13. Wang, L., Zhao, Y., Mueller, K., Kaufman, A.: The magic volume lens: an interactive focus+context technique for volume rendering. In: 16th IEEE Conference on Visualization 2005 (VIS'05), Baltimore, MD (2005) 367–374
14. Bruckner, S., Gröller, M.: Volumeshop: an interactive system for direct volume illustration. In: 16th IEEE Conf. on Visualization 2005 (VIS'05), Baltimore, MD (2005) 671–678
15. Bruckner, S., Gröller, M.: Exploded views for volume data. *IEEE Transactions on Visualization and Computer Graphics* **12** (2006) 1077–1084
16. Ruiz, M., Viola, I., Boada, I., Bruckner, S., Feixas, M., Sbert, M.: Similarity-based exploded views. In: *Smart Graphics*, Rennes, France (2008) 154–165
17. Weiskopf, D., Engel, K., Ertl, T.: Interactive clipping techniques for texture-based volume visualization and volume shading. *IEEE Trans. on Visualization and Computer Graphics* **9** (2003) 298–312
18. Bernhard, O.V., Preim, B., Littmann, A.: Virtual resection with a deformable cutting plane. In: *Smart Graphics*, Magdeburg, Germany (2004) 203–214
19. Chen, H., Samavati, F., Sousa, M.: Gpu-based point radiation for interactive volume sculpting and segmentation. *Visual Computer* **24** (2008) 689–698
20. Wyvill, G., McPheeters, C., Wyvill, B.: Data structure for soft objects. *The Visual Computer* **2** (1986) 227–234
21. Faynshteyn, L.: Context-preserving volumetric data set exploration using a 3D painting metaphor. Master's thesis, Dept. of Computer Science, Ryerson University, Toronto, ON, Canada (2012)