

Layers for Effective Volume Rendering

Sundaresan Raman, Oleg Mishchenko, and Roger Crawfis, *Member, IEEE Computer Society*

Abstract—A multi-layer volume rendering framework is presented. The final image is obtained by compositing a number of renderings, each being represented as a separate layer. This layer-centric framework provides a rich set of 2D operators and interactions, allowing both greater freedom and a more intuitive 2D-based user interaction. We extend the concept of compositing which is traditionally thought of as pertaining to the Porter and Duff compositing operators to a more general and flexible set of functions. In addition to developing new functional compositing operators, the user can control each individual layer's attributes, such as the opacity. They can also easily add or remove them from the composition set, change their order in the composition, and export and import the layers in a format readily utilized in a 2D paint package. This broad space of composition functions allows for a wide variety of effects and we present several in the context of volume rendering, including two-level volume rendering, masking, and magnification. We also discuss the integration of a 3D volume rendering engine with our 2.5D layer compositing engine.

Index Terms—Volume visualization, interaction, volume rendering.

1 INTRODUCTION

Volume rendering has proven to aid in the exploration and understanding of large volumes. While substantial research on transfer functions, volume segmentation, feature detection and rendering integrals has led to impressive results, very little research has gone into user interfaces for volume rendering. This paper builds on the previous research on segmenting and rendering volumes and examines a new system architecture that allows the power and ease of use found in typical 2D paint applications applied to volume rendering.

We segment the workflow process of visualization into two distinct phases. First, traditional volume rendering applied to small regions of interest are rendered. Then, similar to a sort-last scheme for parallel rendering, we apply a compositing phase to the resulting renderings. We expose this compositing phase and provide a rich set of tools and operations to allow unique and effective visualizations that would be difficult in a traditional volume rendering setting. We adopt the concept of layers in a paint package, such as Photoshop [2]. Digital artists and technical illustrators have leveraged the power of the layer concept to quickly provide compositions, to remove or highlight parts of an image, to limit the application of a filter or image adjustment and a variety of additional affects. Often a user will use a copy of a layer to safely experiment with different manipulations or processing of an image.

In this paper, we formalize our concept of layers to the field of volume rendering. We will present several basic examples of processes applied to layers for effective visualization, as well as some more advanced examples that would be difficult to achieve directly in a volume renderer. By allowing layers to be duplicated and processed, a rich set of effects can be achieved without the cost of redundant volume rendering. Additionally, we will present several techniques that mimic the results obtained in recent research reports, but achieved using a much simpler 2.5D paradigm. These techniques are not meant to replace previous work on volume rendering, but rather supplement them. We show examples resulting from our own 3D volume renderer, but the application of layers can be applied to other systems and visualization as a whole.

2 RELATED WORK

There are many well known metaphors for visualization system architectures, from dataflows to spreadsheets. The Application

Programming Environment, or aPe [8], system and the initial Application Visualization System (AVS) [19], combined a dataflow system with a visual programming language. Using a rich set of pre-built data filters and data mappers, the end-user could rapidly construct visualization applications using a computer assisted visual programming system. Even though AVS was intended as a tool for application developers to build turn-key applications for the developer's specific application area, it was adopted by the end-users and used directly for their applications. This was in part due to the clean separation of the visualization process into filters and mappers, the ease of use of the visual programming metaphor and the flexibility to extend the system with user developed modules. This was truly the first visualization system for the masses.

Other visualization systems have mimicked and extended AVS [1][22][16]. IBM Visualization Data Explorer enhances the data flow execution by minimizing redundant computations. The SciRun system allows for rapid prototyping of applications as is evident in their recent PowerApps [17]. Recently, VisTrails [3][6] provides a systematic tracking of the workflow evolution evident in a complex visualization task. By tracking the history or provenance of a visualization session, it allows the user to experiment with possible paths and manipulate the version tree. Compared to other systems, such an approach frees the user from the need for memorizing the changes that led to a particular visualization.

Other unique metaphors for visualization have been few and far between. A spreadsheet approach was developed by Levoy [12] and later extended by Jankun-Kelly and Ma [11]. Levoy uses images and user controls as the values in the cells of the spreadsheet. A general purpose programming language allows complex manipulations with the data. Jankun-Kelly and Ma project the high dimensional visualization parameter space to two-dimensional spreadsheet, and also allow spreadsheet manipulations, such as rotations and translations, bringing additional flexibility to processing the data stored in the spreadsheet.

The specification of transfer functions using genetic algorithms were explored by Marks et al. [14]. With their design galleries approach, the tedious task of tweaking input parameters is done automatically by the system. The system varies the input parameter vector allowing the user to select from a collection of generated results. The system changes the parameter values in a way to make the output well-distributed. The Image Graphs [13] methodology translates the process of data exploration into a graph representation, where each node of the graph consists of an image and the parameters used to produce it. This gives the user a clear representation of the dependencies between the parameter values and the visualization results.

• Sundaresan Raman, Oleg Mishchenko, and Roger Crawfis are with the Department of Computer Science, Ohio State University
E-Mail: {ramansu|mishchen|crawfis}@cse.ohio-state.edu

Our framework exploits the dataflow metaphor. However, we concentrate on the use of layers within the framework. The concept of layers can be seamlessly integrated into all the systems described above as a back-end system.

3 OUR APPROACH

In this section we shall first look at the layers and then at the architecture of the system.

3.1 Layers

The layer concept is central to our framework. By using the term *layer*, we refer to a 2D container for storing an image, on which a number of operations are defined. The intuition for using layers come from the fact that 2D image manipulation is often easier than working with 3D volumes. A number of tasks for volume exploration are performed by processing the layers instead of applying volume rendering techniques.

User interaction with the system combines traditional tasks, but also allows for a rich set of manipulations. Hence, we use the term *layer-centric* to describe the framework. Manipulations include applying a number of operators (programmable shaders) to an individual layer, as well as specifying the order layers are combined. From the user point of view, the way layers are manipulated in our system is similar to the layers' manipulation in Adobe Photoshop. Users can create, delete, change the properties (for example, color and opacity), enable or disable the layer, as well as process the layers with the tools our system provides.

3.2 System Architecture

The central component responsible for the display of the final image is a view. A view is associated with a compositing camera and a viewport. A view also contains a list of layers. Each layer contains a list of drawables, which are the entities that should be rendered. Each drawable has a material associated with it. The material contains attributes needed to render the drawable, for example, color and opacity and a GLSL shader program. The shader provides the implementation of the functionality for the material. All the entities in the hierarchy support one-to-many or many-to-one relationships.

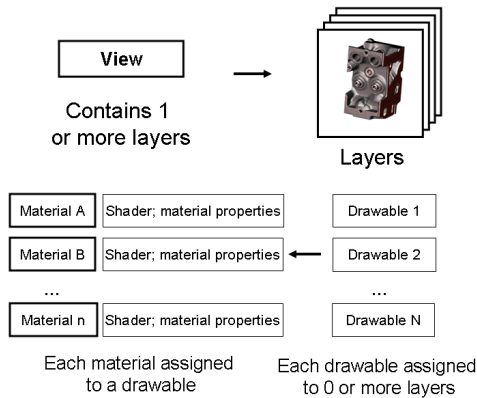


Fig. 1. System architecture

Multiple layers can be rendered with the same drawable or one layer may render multiple drawables. If a layer contains zero drawables, we refer to it as a *null* layer. *Null* layers are generated using shaders which have access to the texture resources in the system. Each layer generates a texture, allowing *null* layers to either create new texture resources or process another layer's texture. The hierarchy of the major components is shown in the Figure 1.

Our system supports transfer function interface with a number of pre-loaded transfer functions. There is also an editor for creating,

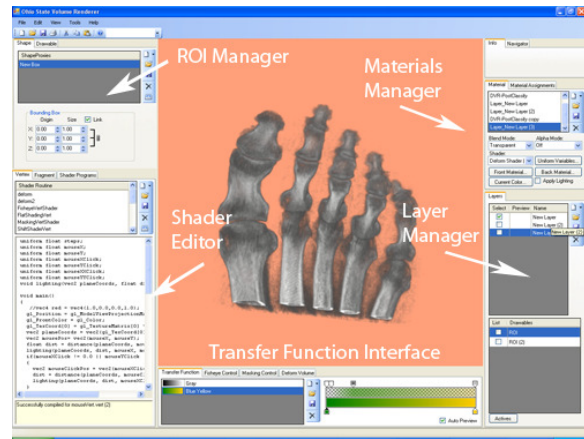


Fig. 2. User Interface Screenshot

editing, compiling and linking shader programs. Compiling and linking are done on the fly.

We employ the layer-based metaphor for the architecture of our visualization system. From the implementation point of view, our framework can be described as a dataflow. However, the difference from other dataflow systems is that we don't explicitly specify the dataflow nor present it to the user. Ease of use in of our framework comes from providing a list of tracked assets and a rich set of user interfaces to build and edit each of the main components: drawables, materials, shaders, layers, and views. Additional assets, such as textures, shared numerical controls and cameras are maintained and tracked by the system for use as shader variables or manipulator controls. Figure 2 presents a screen-shot of our system, showing the shader editor, material edit, 1D transfer-function editor and the layer editor.

The hierarchy of the entities shown in Figure 1 makes it convenient to make the system event based. Whenever any entity changes, the changes are propagated upwards to the entities that are above in the hierarchy, finally leading to the rendering of the image, if necessary. For example, if a material property is changed, it fires an event to all the drawables that share that material. All of these drawables fire events to their parent layers. These layers in turn, fire events to the view and the view fires an event to invalidate the screen.

4 APPLICATIONS OF LAYERS

In this section, we describe some of the applications of the layer based volume rendering. We simulate several recent research contributions using simple 2D operators. While not as general as the 3D contributions, they achieve similar results with lower computational complexity and often a simpler and more intuitive user interface. There are a number of properties that layers provide that are useful to list here before presenting the actual applications. Keeping in mind these properties not only provide better understanding, but also may lead to new ideas of how to use the layers for volume rendering.

We specify the relationships between layers, views and drawables in the following way. Let V be the set of views, L the set of layers and D the set of drawables. We define the function space F as the set of all functions or operators that map a set of drawables to a layer, i.e. $f: D^n \rightarrow L$. Likewise, we define the function space G as the set of all functions that map a set of layers to a view, i.e. $g: L^n \rightarrow V$, as shown in Figure 3. Layer manipulation is specified by the compositing shaders (functions). Each layer is an entity to which functions are applied. We can also define the function space L_F as the set of all functions that map L^n to L . The composition of functions f and g , where f , and g belong to L_F , is denoted as $f \circ g = f(g)$; This operator is not commutative, i.e. $f(g) \neq g(f)$.




Fig. 3. The relationships between layers, drawables and views

A simple set of layer manipulations is the set of image processing algorithms, such as blurring and desaturation. Geometric transformations, such as image warping and distortion, can also be achieved.

4.1 Basic Layer Manipulation

Layers can be seen as a control for interactivity. This is achieved by keeping the set of active layers small. While navigating in a huge dataset, only a low resolution layer may be turned on, providing interactive speed of exploration. Interactivity is especially important for effective volume navigation. By performing volume manipulations, such as rotations, the user often gets a better understanding of the volume. This is due to receiving necessary depth cues, as well as the fact that the human visual system is especially good at detecting movement. When the user zooms in on a selected region of interest, a layer with higher resolution is turned on, providing the necessary level of detail.

Layers are a way to provide the context to the user. Context and regions of interest may be rendered to separate layers and thus processed separately. Such a divide and conquer approach helps to focus on a single layer. This is achieved by turning off the context layers while manipulating the layer in focus. Bringing up the context is done by turning on the inactive layers. On the other hand, in some cases operators that provide *focus + context* functionality may be applied to a single layer.

Layers can be reused. Computationally expensive volume rendering can be avoided in some cases by reusing an already rendered image. When reusing several layers, the overall system performance benefits significantly.

In the next section we provide more concrete examples of applications we have explored with our framework to provide an effective visualization.

4.2 Two Level Volume Rendering

Two level volume rendering [10][9] is an effective technique in volume rendering, combining a number of transfer functions and/or rendering types in the final image. Different transfer functions and/or




Fig. 4. Top: Two Level Volume Rendering with layers. Bottom: The path of a ray through the volume is divided into segments, corresponding to the different materials in the dataset.

rendering modes are better suited for different types of materials in a dataset. In a medical dataset, for example, bones may be rendered with a MIP, while skin and blood vessels with DVR, or vice versa, depending on the user preferences and the particular application. We achieve results similar to the ones of the two level approach by rendering each of the different styles in a separate layer and then compositing them into a final image.

The example is shown in Figure 4. Two different transfer functions are used to render the foot dataset. On the left, DVR is used to create a flat-looking representation of the foot; it is rendered to a top layer. The center image is a DVR rendering of the bones; it is rendered to a bottom layer. In the right is the resulting composited image; the blending level was set to 0.5, i.e. the top layer (skin) is semi-transparent. According to the notation introduced above, there are two functions, DVR_1 and DVR_2 that map D to L . The compositing function maps the resulting two layers to V . This is a rather simple example, using only one drawable and two layers.

The approach described above provides results that slightly differ from the results obtained with the original two level volume rendering. In two level volume rendering, when the ray traverses through the volume, the type of local rendering is selected depending on the type of material encountered in the dataset. The accumulated color and opacity for all ray segments are then composited to produce the final result. Our approach differs in the order in which the segments are composited. In two level volume rendering, the order of compositing corresponds to the order in which the ray encounters the materials. With our approach, all the segments corresponding to the same material in the dataset are composited separately. Then these intermediate results are composited to produce the final image. This is illustrated by the scheme at the bottom of Figure 4. Given two kinds of materials in the dataset, denoted as A and B , the compositing for two level volume rendering gives $I = A_1 * B_1 * A_2 * B_2 * A_3$; while with our approach $I = A_1 * A_2 * A_3 * B_1 * B_2$. Compositing is not commutative, thus in general the results are different, though for some specific cases they may be the same.




Fig. 5. 2D magnification. Lenses applied to both kidneys. Occluding ribs are rendered in a separate layer and are not distorted.

4.3 2D Magnification

Lenses are useful when text or images needs magnification. In volume rendering, 3D lens have been around for about a decade. It is probably the best and simplest example of a focus + context technique. Zooming in a specified region of interest gives the user the necessary detail for exploration, while the region outside the lens keeps the context necessary for volume navigation. This area has received much attention [4][23]. Bier et al [5] came up with a user interface that could employ Magic lens filter to modify the presentation of an object. Viegas et al [20], in 1996, came up with flat lenses and volumetric lenses. Recently, Wang et al [21] developed a framework that provides a variety of GPU-accelerated volumetric lenses.

In our system we use a 2D lens. The zooming effect is achieved by distorting the texture representing the layer. Computationally this allows us to get magnification at almost no cost, since we do not need to re-render the drawables within the layer. Typically, a 2D lens would be applied to the final image. We apply the lens to a single layer (which may be the final image). Applying the zooming

operator to a single layer and not to the final composited image, gives a number of interesting results. First, keeping the context non-zoomed can be useful. An example is shown in Figure 5. The kidney and the ribs are rendered to different layers. The kidneys are focused and magnified, while the ribs are not distorted. We allow for multiple lenses, as shown in Figure 5, where both kidneys are magnified.

There are two ways to achieve 3D magnification; one way is to change the field of view and the other is to traverse through the volume, moving the camera closer to the region of interest. Moving the camera changes the occlusion or ray integration. Thus, 2D lens may be a better choice when magnification is needed without distortion.

4.4 Masking

Masking is a technique used to remove occlusion from a specific area of interest. Given a number of input layers, a masking operator selects and/or blends them according to some specified criteria, such as the opacity. Consider the torso dataset in Figure 9. Keeping the skin visible while trying to explore the skeleton and inner structures, results in a cluttered image. We would like to get an image with the skeleton and other inner structures visible, but only selectively make the skin visible. This can be achieved in the following fashion.

First, skin and skeleton are rendered into two separate layers. The skeleton is rendered to a bottom layer and the skin to the top one. A masking operator makes a part of a top layer transparent, which in turn makes the previous layer visible. An example of this type of masking is shown in Figure 9. The image has two layers composited and the lens makes the enclosed area in the outer layer transparent, thereby displaying the previous layer.

Second, any type of 2D texture could be used as a mask by having a threshold on the opacity values. This approach is used to get the results in Figure 9. The image with skin is rendered on top of the skeleton image. The masking operator checks the opacity value in the skeleton image, and if it is higher than the threshold, blends the skeleton image with the skin layer. Otherwise, only skin is rendered. The user specifies the blending level; in this case, the final image was generated with blending set to 0.8.

There are other flavours of masking functionality that our framework provides. Correa et al [7] came up with a technique for interactive manipulation of volumetric models like deformation or cuts. We simulate volumetric cuts using a 2D cut operator. An example cut of the foot dataset is shown in accompanying video. In Figure 9(e) we show a screenshot from the animation. In Figure 6 we show different styles of masking available to the user. Notice that there may be more than one location to apply a masking operator to as well as it is possible to apply the operator to different layers.




Fig. 6. Different masking styles.

4.5 Slab-based Rendering

Layers are most useful when the data is segmented into several parts and each part is rendered to a separate layer. We have had much success in applying a two-step volume rendering based on slabs. A volume is partitioned into slabs that are either view-aligned or aligned to the world axis most perpendicular to the desired viewing plane. Each slab is rendered to a corresponding layer. In the accompanying animated gif, we use view aligned slabs and show one of the applications of slabs: slab-based traversal through the volume. Slab-based traversal can be viewed as an application of a meta-operator to a set of layers. The meta-operator specifies the operator that should be applied to a particular layer (and thus slab), as well as the time when the operator should be applied. A simple example is an application of a window function to a set of layers. We implement the traversal by changing the opacity of the current layer, starting from the top one.

The results of slab-based traversal are similar to the results obtained by the use of a slice plane that is moved through the volume. However, our slab-based approach has a major advantage. Every change in the location of the slice plane results in re-rendering of the volume. With slabs, once they are rendered, traversal is based only on changing opacities of the corresponding layers.

In its simplest form, this can be thought of as a voxelization and volume rendering. The voxelization process integrates rays through the slab and stores the result to a layer. The final compositing is a simple volume rendering with a small number of slices. We typically use less than ten slabs. As such, very fast and efficient compositions are possible. Of course, a much broader set of compositing operations are also possible.

4.6 Drop Shadows

Figure 7 shows a compositing operation that produces a drop shadow. The volume is rendered to a layer and a copy of the layer is made. The shadow operator processes the layer and colours the parts with non-zero opacity with black. A shift operator moves the rendering of the shadow. Finally, the volume rendering and the shadow renderings are composited to get the result shown in the figure. This example illustrates an




Fig. 7. Drop shadow.

important concept of layer reuse. We could render the volume twice with a different transfer function to achieve the same result; however, our layer approach is superior in terms of computational cost.

4.7 Blur and Desaturation

Keeping the amount of visual information presented to the user limited is crucial for effective volume exploration. Too many details can distract the user from the features he/she is interested in. One effective way of dealing with the problem is defocusing or desaturating the context, keeping high-level detail and/or colors only in the region of focus. Another option is to use a masking pattern, such as hatch pattern. Examples of these approaches are shown in Figure 8.



Fig. 8. Left: Gaussian blur. Center: Desaturation. Right: Hatch pattern.

This is an example of layer reuse. Instead of utilizing two transfer functions and re-rendering the volume twice, by processing a layer or a set of layers we can achieve the desired effect. (in the above example a single layer keeps both the context and the region of interest). This example along with the previous one illustrate the flexibility of layers: if necessary, multiple layers can be used to achieve the desired result; however, for some applications a relatively simple operator and only one layer are enough.

4.8 Creating Scientific Illustrations

One interesting application that our framework is suitable for is creating scientific illustrations. An example is shown in Figure 9(g). To make this image, we used two features of our framework: exporting and importing the layers. Grid texture was loaded in a separate layer and the head rendered on top of it with opacity set to 0.8. The resulting image was exported to Photoshop, where the labels were added.

Texture import can be useful for loading any kind of ruler or some other background. This is a common practice in digital art, where the background may be a hand-drawn sketch to use as a template. It will not be used in the final composition, but provides a valuable reference. Likewise, during a volume rendering session, it may be useful to display an image of a reference model, allowing the user to determine the settings and shaders to best match the reference point of view and feature highlights.

5 IMPLEMENTATION

Our system is implemented in C# and OpenGL [15], with the use of the Tao framework [17] that provides OpenGL bindings for C#. Shader programs are written in GLSL. We use 3D texture-based volume rendering with back to front compositing for the attached volume renderer.

The system is event based. Any change in any of the entities in the architecture hierarchy forces the events to fire and propagate through the hierarchy. The components that are registered to receive a particular type of event perform the necessary actions to update their state upon receiving the corresponding event. This keeps the number of state changes minimal. For example, if user changes a transfer function for one specific region of interest, only events corresponding to this change are fired. This finally leads to re-rendering of the specified region of interest. No other components are updated and unnecessary redraws are not performed.

While performance was not a concern when designing the system, we achieve a number of performance benefits. Eliminating unnecessary redraws as described above can bring significant performance improvements. This is especially noticeable with multiple ROIs and layers. Other performance advantages come directly from the layer properties and the way layers are utilized for particular tasks. As mentioned previously, we can control interactivity by turning off some expensive volume renderings while volume exploration. This is most applicable whenever there is a camera change event that makes all the layers invalid. Also, reuse of layers, as we saw in section 4, helps us avoid redundant volume renderings.

6 CONCLUSION

We have presented a layer-based volume rendering system. The ease of use and flexibility of layers allow the user to perform a variety of volume exploration tasks. In our future research work, we plan to automate the process of layer manipulation. With the current implementation of the system, the user is given full control over of layer manipulation, including layer creating. This, however, implies that the user should carefully consider how to select layers and operators for a particular task. Often this is a tedious and trial-and-error process. In case the layers, regions of interest and operators are not selected properly, the result may be of no use. Thus our goal is to make the process of translating the problem into layers paradigm automated or semi-automated. For example, the user may be suggested to select a number of volume exploration scenarios when loading datasets into the system.

We are also going to focus on further development of the system architecture. Current design allows separating the system into independent components, making the system suitable for remote volume rendering applications and collaborative visualization. Datasets may be rendered at a powerful graphics workstation and the results may be sent over the network to other computers or handheld devices. This also brings the questions of the user interfaces on the smaller screens, as well as the questions of how to distribute the most computationally expensive jobs in case there is more than one powerful graphics machine.

REFERENCES

- [1] G. Abrams and L. Trenish, "An Extended Data-Flow Architecture for Data Analysis and Visualization," Proc. IEEE Visualization 1995, IEEE CS Press, Los Alamitos, Calif., 1995, pp. 263-270.
- [2] Adobe Photoshop. <http://www.adobe.com/>
- [3] L. Bavoil, S.P. Callahan, P.J. Crossno, J. Freire, C. E. Scheidegger, C.T. Silva, and H. T. Vo. VisTrails: enabling interactive multiple-view visualizations In *Proceedings of IEEE Visualization*, 2005, Vis'05, 135-142
- [4] E. Bier, M.C. Stone, K. Pier, W. Buxton, and T.D. DeRose, 1993. Toolglass and Magic Lenses: The Seethrough Interface, *Proc ACM SIGGRAPH 1993*, Anaheim, CA, Aug 1993, 73-80.
- [5] E. Bier, M.C. Stone, and K. Pier, 1997. Enhanced Illustration Using Magic Lens Filters. *IEEE Comput. Graph. Appl.* 17, 6 (Nov. 1997), 62-70.
- [6] S. Callahan, J. Freire, E. Santos, C. Scheidegger, C. Silva, and H. Vo, VisTrails: Visualization Meets Data Management. *SIGMOD 2006*, June 27-29, 2006, Chicago, IL
- [7] C. D. Correa, D. Silver, and M. Chen. Feature Aligned Volume Manipulation for Illustration and Visualization. *IEEE Transactions on Visualization and Computer Graphics* 12, 5 September 2006
- [8] Dyer, D. S. 1990. Visualization: A Dataflow Toolkit for Visualization. *IEEE Comput. Graph. Appl.* 10, 4 (Jul. 1990), 60-69.
- [9] M. Hadwiger, C. Berger, and H. Hauser, 2003. High-Quality Two-Level Volume Rendering of Segmented Data Sets on Consumer Graphics Hardware. In *Proceedings of the 14th IEEE Visualization 2003 (Vis'03)* (October 22 - 24, 2003). IEEE Visualization. IEEE Computer Society, Washington, DC, 40.
- [10] H. Hauser, L. Mroz, G. Bischl, and M. Gröller. 2001. Two-Level Volume Rendering. *IEEE Transactions on Visualization and Computer Graphics* 7, 3 (Jul. 2001), 242-252.
- [11] T.J. Jankun-Kelly and K.-L. Ma, "A Spreadsheet Interface for Visualization Exploration," In Proc. Visualization 2000 Conf., ACM Press, New York, Oct. 2000.
- [12] M. Levoy, 1994. Spreadsheets for images. In *Proceedings of the 21st Annual Conference on Computer Graphics and interactive Techniques SIGGRAPH '94*. ACM Press, New York, NY, 139-146.
- [13] K.L. Ma, "Image Graphs: A Novel Approach to Visual Data Exploration, *Proc. IEEE Visualization '99*, pp. 81-88, 1999.
- [14] J. Marks, B. Andalman, P. Beardsley, W. Freeman, S. Gibson, J. Hodgins, T. Kang, B. Mirtich, H. Pfister, W. Rumli, K. Ryall, J. Seims, and S. Shieber, 1997. Design galleries: a general approach to setting parameters for




Fig. 9. Top Row: (a) skeleton and (b) skin rendered to separate layers. (c) The result of applying masking operator to (b) and (a). (d) Without masking, image is cluttered. Bottom Row: (e) Volumetric cut. (f) Masking with lens. (g) Scientific Illustration made with our system.

computer graphics and animation. In *Proceedings of the 24th Annual Conference on Computer Graphics and interactive Techniques International Conference on Computer Graphics and Interactive Techniques*. ACM Press/Addison-Wesley Publishing Co., New York, NY, 389-400.

- [15] OpenGL. <http://www.opengl.org>
- [16] S.G. Parker and C.R. Johnson, 1995. SCIRun: A scientific programming environment for computational steering. In *Proceedings of the 1995 ACM/IEEE Conference on Supercomputing (Cdrom) – volume 00* (San Diego, CA, December 04-08, 1995) Supercomputing '95. ACM Press, New York, NY, 52.
- [17] PowerApps. <http://software.sci.utah.edu/scirun.html>
- [18] Tao Framework, <http://www.taoframework.com>
- [19] C. Upson, T. Faulhaber, D. Kamins, D. H. Laidlaw, D. Schlegel, J. Vroom, R. Gurwitz, and A. van Dam, 1989. *The Application Visualization System: A Computational Environment for Scientific Visualization*. *IEEE Comput. Graph. Appl.* 9, 4 (Jul. 1989), 30-42.
- [20] J. Viega, M. Conway, G. Williams, and R. Pausch, 1996. 3D Magic Lenses. In *Proceedings of the 9th Annual ACM Symposium on User Interface Software and Technology*. Seattle, WA, Nov 06-08, 1996, UIST 96, ACM Press, New York, NY, 51-58.
- [21] L. Wang, Y. Zhao, K. Mueller, and A. Kaufman, 2005 The Magic Volume Lens: An Interactive Focus+Context Technique for Volume Rendering. In *Proceedings of IEEE Visualization '05*, 367-374, 2005.
- [22] M. Young, D. Argiro, and S. Kubica, 1995. Cantata: visual programming environment for the Khoros system. *SIGGRAPH Comput. Graph.* 29, 2 (May. 1995), 22-24.
- [23] J. Zhou, M. Hinz, and K. D. Tonnes. Focal region-guided feature-based volume rendering. In *Proceedings 1st International Symposium on 3D Data Processing Visualization and Transmission (3DPVT)*, pages 87--90, Padova, Italy, June 2002.