

State of the Art in Transfer Functions for Direct Volume Rendering

Patric Ljung¹ Jens Krüger^{2,3} Eduard Gröller^{4,5} Markus Hadwiger⁶ Charles D. Hansen³ Anders Ynnerman¹¹Linköping University, Sweden ²CoViDAG, University of Duisburg-Essen³Scientific Computing and Imaging Institute, University of Utah ⁴TU Wien, Austria ⁵University of Bergen, Norway⁶King Abdullah University of Science and Technology

Abstract

A central topic in scientific visualization is the transfer function (TF) for volume rendering. The TF serves a fundamental role in translating scalar and multivariate data into color and opacity to express and reveal the relevant features present in the data studied. Beyond this core functionality, TFs also serve as a tool for encoding and utilizing domain knowledge and as an expression for visual design of material appearances. TFs also enable interactive volumetric exploration of complex data. The purpose of this state-of-the-art report (STAR) is to provide an overview of research into the various aspects of TFs, which lead to interpretation of the underlying data through the use of meaningful visual representations. The STAR classifies TF research into the following aspects: dimensionality, derived attributes, aggregated attributes, rendering aspects, automation, and user interfaces. The STAR concludes with some interesting research challenges that form the basis of an agenda for the development of next generation TF tools and methodologies.

Categories and Subject Descriptors (according to ACM CCS): I.3.8 [Computer Graphics]: Applications—Volume Rendering I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture I.4.10 [Computer Graphics]: Image Representation—Volumetric

1. Introduction

A transfer function (TF) maps volumetric data to optical properties and is part of the traditional visualization pipeline: data acquisition, processing, visual mapping, and rendering. Volumetric data is considered to be a scalar function from a three-dimensional spatial domain with a one-dimensional range (e.g., density, flow magnitude, etc.). Image generation involves mapping data samples through the TF, where they are given optical properties such as color and opacity, and compositing them into the image.

A TF simultaneously defines (1) which parts of the data are essential to depict and (2) how to depict these, often small, portions of the volumetric data. Considering the first step, a TF is a special, but important, case of a segmentation or classification. With classification, certain regions in a three-dimensional domain are identified to belong to the same material, such as bone, vessel, or soft tissue, in medical imaging. A plethora of classification and segmentation algorithms have been developed over the last decades, with semi-automatic approaches often tailored to specific application scenarios. Segmentation algorithms can be quite intricate since information from the three-dimensional spatial domain and the one-dimensional data range are taken into account. TFs in their basic form are, on the other hand, restricted to using only the data ranges. In comparison with general classification algorithms, this characteristic makes a

TF less powerful with respect to identifying relevant parts of the data. The advantage of a TF is, however, a substantial performance gain as classification based on the one-dimensional data range reduces the complexity tremendously. This gain is the result of the three-dimensional domain being typically two orders of magnitude larger than the small data range. Histograms are an example of discarding potentially complex spatial information and aggregating only the data values to binned-frequency information. In the same spirit, TFs classify interesting parts of the data by considering data values alone. The second functionality of a TF deals with specifying optical properties for portions of the data range previously identified as being relevant.

A survey of the works published in the field of TFs reveals that in the three decades since the earliest published techniques for direct volume rendering [KVH84, DCH88, Lev88], over a hundred studies have been published in the most influential journals and conferences. In 2001, the Transfer Function Bake-Off examined four of the then most promising approaches to TF design [PLB*01]: trial and error with a TF editor; data-centric using computed metrics over the scalar field; data-centric using a material boundary model; and image-centric using an organized sampling of exemplar images. In 2010, Arens and Domik [AD10] authored a survey of TFs for volume rendering in which they subdivided TFs into the follow-

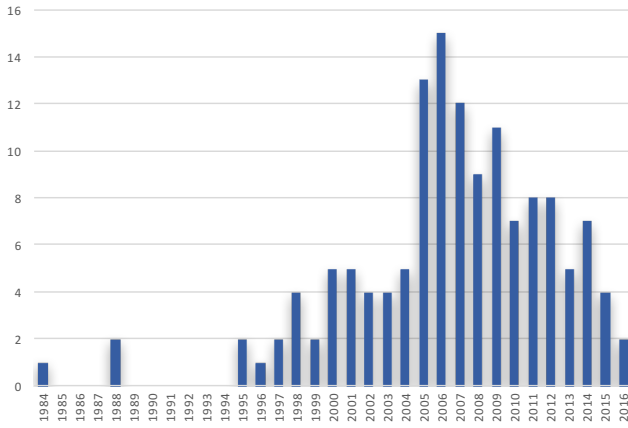


Figure 1: Number of publications related to transfer functions over the last three decades.

ing six categories: 1D data-based, gradient 2D, curvature-based, size-based, texture-based, and distance-based. These six categories were compared using automaticity, user-interaction, constraints, emphasis, memory consumption, and real-time capability. The survey concluded that no single TF is generally applicable for all situations, and expertise is required to determine the most suitable TF for specific data and/or application.

During the six years since the last published survey, significant progress on TFs has been made. Figure 1 provides a histogram, not claimed to be complete, showing the numbers of papers published on the topic of TFs during the past 32 years, including several since the last survey in 2010. The goal of this state-of-the-art report (STAR) is to fill the gap in the literature review since the last survey by bringing together a diverse group of researchers in the field of direct volume rendering. We provide an up-to-date overview of the entire area of TFs for direct volume rendering, from the earliest approaches to the most recent advances.

1.1. Structure and categorization

This STAR follows the categorization shown in Figure 2. Section 2 reviews the fundamentals of TFs. Section 3 examines TFs from the dimensionality and native data domain perspective. TFs can include multiple derived attributes, discussed in Section 4, or aggregated attributes, described in Section 5. These sections on data domain aspects relate mostly to the *material classification* component of the TF. Section 6 reviews research about how TFs can be used in the rendering process to focus attention on relevant data, referred to as *visual mapping*. Two usability aspects of the TF are covered in the following sections. Methods and techniques for aiding the user in TF settings, via levels of automation, are presented in Section 7. User interfaces, and the topics on how the user can interact with a system to define a particular TF setting, are described in Section 8. Some potential areas of future research are presented in Section 9, while Section 10 concludes the STAR.

This STAR aims to provide a comprehensive account of trends and approaches spanning the TF landscape. The STAR is also in-

tended to show that the TF maintains its instrumental role in visual data analysis by introducing novel and effective approaches to TF design and use.

2. Transfer Function Fundamentals

In this STAR, we assume that the reader has an elementary understanding of the basic concepts involved in direct volume rendering (DVR). For a general overview and introduction to volume rendering, we refer the reader to Engel *et al.* [EHK*06], in particular the sections on TFs in Chapters 4 and 10. This background section provides the basic concepts defining the role of the TF and the complexity that it abstracts away. The section also problematizes the multiple roles of the TF as a material classifier and carrier of optical properties. An overview of the evolving and expanding role of the TF in the wider context of the DVR pipeline is also provided.

The starting point for our overview is the well known emission-absorption volume rendering equation

$$\mathbf{I} = \int_a^b \mathbf{q}(s) e^{-\int_a^s \kappa(u) du} ds, \quad (1)$$

where \mathbf{I} is the resulting light intensity after traversing the ray between the points a and b in the volume, and $\mathbf{q}(s)$ is the light contribution from the point s along the ray. The attenuation of the light contribution along the ray is estimated using the optical depth $\tau = \int \kappa$. The functions $\mathbf{q}(s)$ and $\tau(u)$ are defined by the properties of the material(s) in the object of study and the transport of light through the material(s). The role of TFs in volume rendering is to provide estimates of these two functions. It is important to realize that the seemingly simple mapping from data to optical properties hides layers of complexity. In its most rigorous interpretation, the TF would be based on the simulation of physical light transport in participating media containing all the materials present in the object of study. However, in practice, this is not possible due to computational limitations nor is it desirable in most applications. The data may not represent a physical object, and/or optical properties corresponding to a certain data value are based on user design principles rather than on a representation of the physical transport of light. Furthermore, TFs are exploration tools that are interactively changed during volume exploration to emphasize regions of interest. For more details on light transport in participating volumetric media, we refer readers to the seminal paper by Max [Max95], in which optical models for DVR are discussed.

In its discrete version, Equation 1 can be viewed as a compositing operation over color contributions from sampled points along a ray. With one such ray per pixel, originating in the image plane, the final pixel color is determined by the forward compositing equation

$$\mathbf{I} = \sum_{i=1}^n \mathbf{C}_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j), \quad (2)$$

where \mathbf{C}_i is the color and α_i is the fraction of incoming light that is absorbed at the discrete position i along the ray. This equation can be solved front-to-back using the recursive forms

$$\mathbf{C}'_{i+1} = \mathbf{C}'_i + (1 - \alpha'_i) \mathbf{C}_i \alpha_i \quad (3)$$

$$\alpha'_{i+1} = \alpha'_i + (1 - \alpha'_i) \alpha_i \quad (4)$$

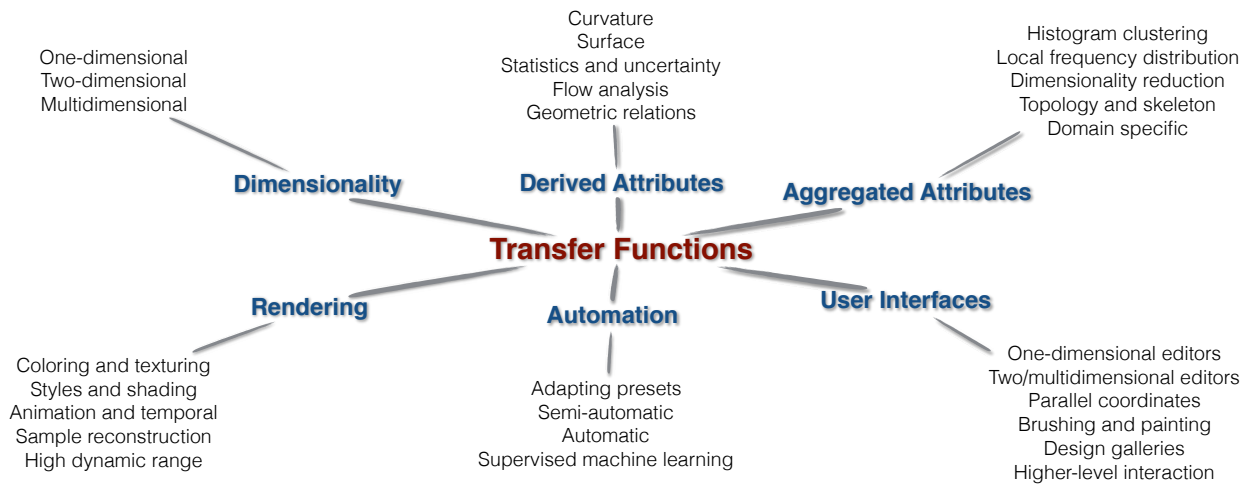


Figure 2: The transfer function classification used in this STAR.

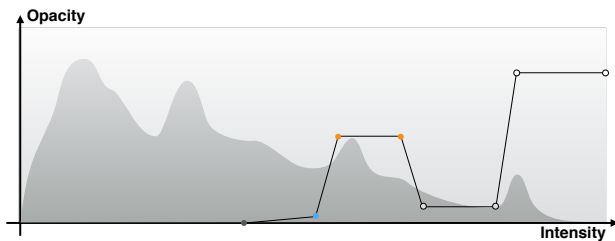


Figure 3: A standard user interface for 1D TF settings. Color and opacity are assigned to data ranges using piecewise linear TF widgets. The background represents the histogram of binned scalar attribute data.

where \mathbf{C}'_i is the opacity-weighted accumulated color-contribution of ray samples 1 up to i , and α'_i is the accumulated opacity, with initial color and opacity $\mathbf{C}'_1 = 0$, $\alpha'_1 = 0$.

Here the role of the TF is to perform a mapping of data values to optical properties, resulting in estimates of \mathbf{C}_i and α_i . Figure 3 shows a standard application of a TF to a data volume with a single scalar. The histogram of the data serves as the basis for the 1D TF widgets that are used to assign color and opacity to data ranges.

It is interesting to note that initial publications on volume rendering consider a TF conceptually but do not use the explicit term. Kajiya and Von Herzen [KVH84] were among the first to introduce the concept of volume rendering of a 3D density field. They provided an approximation of the full radiative scattering and, furthermore, showed that a simple forward scattering approach is insufficient to render phenomena such as clouds. Levoy [Lev88] presented the first paper using the gradient magnitude to enhance the boundaries, or rather suppress the interior of homogeneous regions in the volume. Although Levoy did not use the term TF, the mapping to colors and opacities was described. In the same year, Drebin *et al.* [DCH88]

described the mapping from data values to optical properties using a probabilistic view of material presence and mixture.

2.1. The dual role of the transfer function

As an illustration of the TF complexity, let us assume that the object of study consists of a number of materials with different optical properties. The TF is seen as a material classification tool in both perspectives, but there is a difference in whether material probabilities are an *implicit* or *explicit* part of the mapping. In the explicit case, the application of a TF is modeled as a two-step approach. First, the sample value s is mapped to a set of material probabilities $p_m(s)$, where m is the index among the M materials. Then, the material probabilities are used to combine the individual material colors $\mathbf{C}_m = (r_m, g_m, b_m, \alpha_m)^T$, which results in the sample color \mathbf{C} . Such an approach was employed in the initial DVR implementation by Drebin *et al.* [DCH88] and further thoroughly elaborated upon by Kniss *et al.* [KUS*05] while studying the probabilistic classification, and later by Lundström *et al.* [LLPY07], in the context of uncertainty visualization. In the implicitly probabilistic view, the TF is seen as a direct mapping from sample value s to sample color \mathbf{C} . This view is currently the dominant approach, and it is the view represented in recent DVR literature [EHK*06].

2.2. The evolving landscape of the transfer function

The research compiled in this report demonstrates how the TF concept has evolved into increasingly advanced TF designs and uses. As shown in Figure 4, the landscape in which the TF now resides goes far beyond its traditional domain in which the TF acts only as classifier and mapper of material properties.

Data that needs to be visualized is becoming increasingly multi-modal in nature, which enables improved classification capabilities. For example, modalities are combined to produce co-registered data fields in the area of medicine. Relevant examples are found in the increasing use of CT-PET and MR-PET, but also in the development

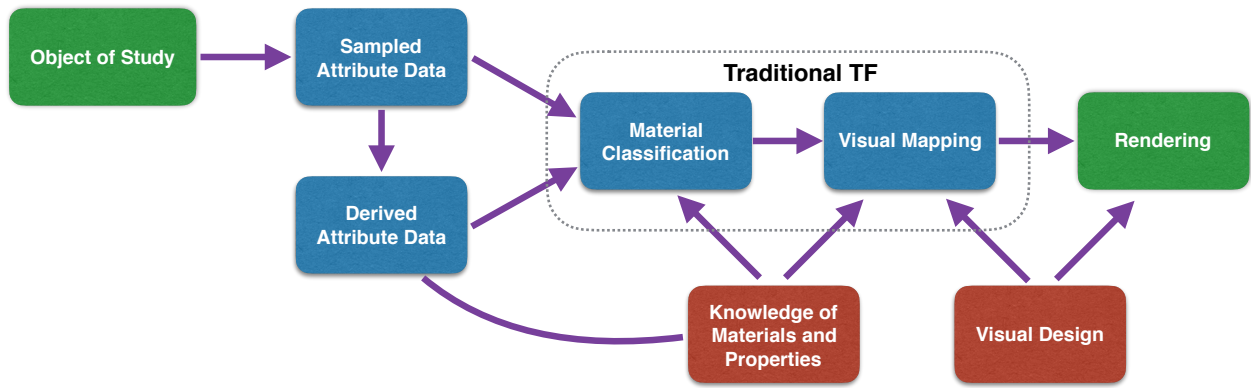


Figure 4: Embedding of the TF into object analysis, visual design, and image generation through rendering.

of multispectral CT, which will significantly increase the dimensionality of both sampled and derived attributes. The rapid move towards multimodality calls for TF approaches that can deal with high-dimensional attribute data. Multidimensional data can be uncorrelated, which opens up the possibility of using separable 1D TFs that are later combined to form complete multidimensional TFs.

Another trend is the use of a TF as a tool for the expression of user knowledge of material presence and properties, which transforms the TF into a source of information that can be used to derive further attribute data. High level expressions of knowledge codify the complex relations between data domains and the corresponding TFs and indeed between different TF segments. A powerful scheme is to define the TF as material presence components, which follows the ideas presented by Lindholm *et al.* [LLL*10], where labels are assigned to TFs with semantics, indicating, for instance, bone, muscle, fat, etc. Several important contributions to the field rely on the expression of knowledge encoded in the TF in ways that go beyond the traditional TF domain. Such approaches have been used, for instance, in data reduction and multiresolution representations.

Perhaps one of the most important aspects of the TF is the freedom it provides users in designing visual appearance. The TF further allows users to interactively alter the visual parameters and their application to different parts of the object of study. The number of attributes is increasing, and the usability of available TF tools is improving. Thus, the TF is increasingly becoming an instrument for the exploration of scientific content of the data to produce, both scientifically relevant and aesthetically appealing, high quality images. The role of the TF as a design and workflow tool was already apparent in the Transfer Function Bake-Off [PLB*01], which analyzed four approaches in the TF design and highlighted differences in the workflow of the TF design process.

3. Dimensionality

Volumetric data is commonly presented as scalar, bi-variate, or multivariate. Examples of scalar data are medical data from CT scanners and particle density fields from numerical simulations. Bi-variate data is also found in medical imaging from dual energy CT scanners

as well as in complex fields, with real and imaginary values, from numerical simulations in many different domains. There are many examples of multivariate volumetric data, often from numerical simulations in fluid dynamics, that create vector or tensor data, but also from seismic surveys and astrophysical numerical simulations. Multimodal data is also a common form, in which the data may be represented on grids with different extent and alignment, such as Ultrasound B-mode data together with CT data.

As the starting point for investigations in most application domains is data, we initially examine the data perspective. The data on which the TF is operating defines the corresponding *data domain*, or simply the *transfer function domain*, i.e., the domain in which the function is defined, and the corresponding dimensionality. We sub-categorize this section in terms of dimensionality, which covers both the source dimensionality as well as simpler derived properties. Later we describe the subcategory of first- and second-order types of additional local attributes used for the TF, such as gradient magnitude, curvature, etc. In the following subsections, we survey the literature on 1D, 2D, and multidimensional (MD) TFs and focus on the material classification aspect of the TF, rather than the visual mapping.

3.1. One-dimensional data

The simplest kind of TF has a one-dimensional domain where the function is defined, i.e., it is a 1D TF operating on a scalar input value. This input is most commonly the scalar value given by the scalar field comprising the input volume, such as the material density. The one-dimensional TF classifies the scalar data value, d , and subsequently maps the material to an optical property for rendering:

$$\mathbf{q}(d) = \mathbf{V}(\mathbf{M}(d)),$$

where $\mathbf{M}(\cdot)$ is the material classification function and $\mathbf{V}(\cdot)$ is the visual mapping of the material.

In the paper by Drebin *et al.* [DCH88], the concept of the TF is presented as a material classification with probabilities based on the scalar value, and thereafter a visual mapping is applied. The authors point out that material classification is a probabilistic and

not a binary decision, leading to the notion of material mixtures and how to blend or mix the visual mappings in the presence of multiple materials. See Section 2.1 for a discussion on the dual role of the TF and Section 6 for visual mapping strategies.

Without additional attributes, the 1D TF is quite straightforward, and in a strict interpretation, rare, if we consider the visual mapping to be a part of the TF. Visual mapping often requires the normal for shading operations and thus the normal constitutes an additional input. One example of a strict 1D TF is the Gradient-Free Shading described by Desgranges *et al.* [DEP05]. The noisy nature of 3D Ultrasound data is a good example of a case in which 1D TFs are suitable. However, it may be beneficial to include ray depth to provide improved depth-cues as performed in the work of Srinivasan *et al.* [SLMSC13]. Nevertheless, this issue is more of a visual mapping aspect rather than one of classification.

One-dimensional TFs are adequate in many cases of simulation data where measurement noise is low or even non-existent. Other examples include industrial CT scans, such as in Li *et al.* [LZY*07b], where different materials of interest have few overlapping intensity ranges. For medical image data, the 1D TF is often inadequate as tissues have significant overlap in the intensity range, as described by Lundström *et al.* [LLY06a]. In addition, medical data is measured and relatively noisy, which further negatively impacts the ability of 1D TFs to correctly classify different tissue types.

Despite their shortcomings, 1D TFs are the most common form of TFs, especially outside the visualization research community. One-dimensional TFs are often the first tool available in software packages providing volume rendering, as they are relatively easy to comprehend for the novice or occasional user. Practically all production visualization software, such as ParaView [HA04], VisIt [CBW*12], or ImageVis3D [CIB15], support 1D TF editors. Constructing a 1D TF is most commonly achieved by combining separate TF components, which simplifies several aspects, such as user interface interactions and adaptation to new datasets, as discussed in Castro *et al.* [CKLG98].

3.2. Two-dimensional data

If the TF operates on an input with more than one dimension, it is termed a 2D TF (for a bi-variate input) or an MD TF (for an input of multiple dimensions).

A distinction needs to be made whether the TFs are separable or if they are intrinsically high-dimensional. A separable 2D TF is defined as two separate 1D functions that are combined only after both 1D functions have been applied separately, which is most commonly done via the tensor product. For example, the first dimension may define the material classification, and the second dimension controls some aspect of the visual mapping. Non-separable 2D TFs are better able to classify materials or features in the data compared to separable 2D TFs, which are essentially two separate functions combined afterward:

$$\mathbf{q}_{\text{non-separable}}(d_1, d_2) = \mathbf{V}(\mathbf{M}(d_1, d_2)),$$

and

$$\mathbf{q}_{\text{separable}}(d_1, d_2) = \mathbf{V}(\mathbf{M}(d_1), d_2).$$

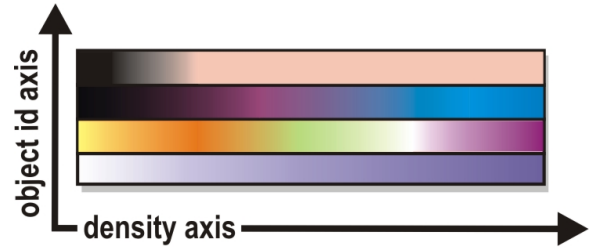


Figure 5: 2D TF texture combining material density and object (label) ID of segmented objects [HBH03]. Image courtesy Hadwiger. Copyright 2003 IEEE.

The latter formulation defines a separable 2D TF, with examples such as gradient-based opacity modulation in which the second dimension is used to improve visual appearance, which suppresses interior homogeneous material regions and enhances boundaries. The material classification power of a separable TF still corresponds to that of a 1D TF, but, by also including higher dimensions, provides a significant enhancement to the visual appearance. Essentially, a 1D TF is applied first, followed by multiplying the opacity with a 1D function of gradient magnitude. This type of 2D TF is very easy to store, because it can be represented as two 1D functions instead of as a full 2D function. In the case of gradient magnitude-weighted opacity, the second function is described by a simple equation and does not need to be stored as a lookup table. A classic example is the gradient magnitude-weighted opacity modulation of the classified value. The seminal work by Levoy [Lev88] contains an example of this type of 2D TF, but the approach cannot be considered as fully separable.

A very simple, yet important, type of non-separable 2D TF is the use of multiple 1D TFs for rendering segmented volume data. See Figure 5 for an example of such a TF for a volume containing four segmented objects. In this case, one dimension is the scalar value, and the second dimension is simply the ID of a segmented object (also called a label ID). We note that this type of 2D TF is not separable, but it also does not really constitute a “general” 2D TF, because multiple 1D TFs are simply combined into a 2D table in a straightforward manner. However, this approach, or a variant thereof, is used in many volume renderers that render segmented data [HBH03, BG05]. The basic approach can be extended further, for example in style TFs for segmented volumes [BG07] (see Figure 11). In essence, this type of 2D TFs for segmented volumes delegates (most of) the material classification to the segmentation process that occurs prior to the visualization stage. Segmentation augments the unsegmented volume dataset with a label ID for each voxel, which can then be directly fed as the second input into this type of TF.

An example of a genuinely non-separable 2D TF is the value and gradient magnitude mapping presented in Kniss *et al.* [KKH02], which cannot be obtained as the tensor product of two 1D TFs. In this work, 2D TF widgets (polygons) are defined using a graphical user interface over the desired region in the 2D frequency distribution plot, or 2D histogram (see Figure 6). It should be noted that in contrast to a separable 2D TF, a non-separable 2D TF requires storing (or representing) a full 2D function, which means that it

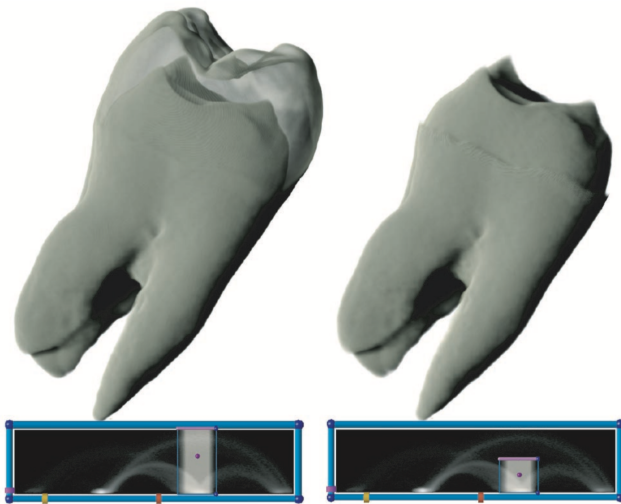


Figure 6: The tooth dataset in which a general non-separable 2D TF (density vs. gradient magnitude) is used to more precisely select features in the data [KKH02]. Image courtesy of Kniss. Copyright 2002 IEEE.

usually requires storing a 2D array or a 2D texture. Both approaches to non-separable TFs can also be combined in a straightforward manner, i.e., using one general 2D TF per segmented object, which amounts to storing all 2D TFs in a higher-dimensional table.

3.3. Multidimensional data

If the user is to manipulate the TF definition directly, moving beyond 2D TFs immediately poses significant challenges in terms of user interfaces and cognitive comprehension. Much research and work on MD TFs is, therefore, related to various forms of automation and advanced user interfaces, which are covered in Sections 7 and 8. Typical approaches include dimensional reduction, clustering and grouping, machine learning, and various user interface approaches such as parallel coordinates or direct slice or volume view interaction.

From a material classification point of view, several methods investigate different features, which include distributions around a point, such as radial basis functions (RBFs), or various geometric primitives, such as box, pyramid, or ellipsoid primitives. Kniss *et al.* [KKH02] employ box and pyramid shapes, but only two dimensions are shown at a time. Alper, Selver and Guzelis [SG09] use RBFs in the classification scheme, discussed in Section 7.2.

Methods using separable classifiers allow for easier definitions but provide a weaker classification power in the MD data space. Zhou *et al.* [ZSH12] discuss combinations of 2D primitives and 1D TFs. Here each 2D primitive has its own associated 1D TF where the data domain can be chosen by the user. Using separate 1D TF definitions for each widget, or region, defined in the 2D domain improves the classification power over using a single shared 1D TF. Additional user interface aspects are covered in Section 8.2.

Classification spaces with related attributes, such as data value range, range or distribution over first- and second-order derivatives, sometimes allow parametrization of some dimensions. This can make the TF definition easier [HPB*10].

Industrial spectral X-ray CT data has been tackled by Amirkhanov *et al.* [AFK*14], who define MD TFs directly for the spectral data to identify material compositions in industrial components. The system also includes X-ray fluorescence spectral data. A color-mix is computed based on the spectral energy of a spectral bin and its associated color, and there is no explicit classification of materials.

Another form of MD TFs deals with the classification of color data, as in the cases of the Visible Human Project, which generated slice stacks of photographs. The challenge here is to produce a useful volume rendering that reveals structures within the data, since color ranges are highly overlapping and ambiguous. Morris and Ebert [ME02] propose a method to map color data to opacity TFs to reveal the interior structures of the data. They use boundary enhancement techniques, which incorporate the first and second directional derivatives along the gradient direction. The CIEL*u*v* color space is chosen in their algorithms, and a deeper analysis is provided in Ebert *et al.* [EMRY02]. For medical RGB color data, Takanashi *et al.* [TLMM02] perform independent component analysis (ICA) of the RGB-color histogram space. Their assumption is that moving a clipping plane along the ICA axes allows the user to discriminate different parts of the data. Muraki *et al.* [MNKT01] use a similar ICA-based approach to transfer colors from RGB datasets to multichannel MRI volumes.

4. Derived Attributes

So far, we have described what could be considered the “traditional” attributes of volume data that are used as the input to TFs in the form of *scalar value* (density), *gradient magnitude*, and *object/label ID* (in the case of segmented data). In this section, we consider additional derived attributes that are often used.

4.1. Curvature

The curvature of surfaces is an important attribute that characterizes their local *shape*. Even for volume rendering, the computation and use of curvature makes sense. Usually, employing this attribute means interpreting curvature as *isosurface curvature* at a specific position, which does not necessarily require computing an actual isosurface. The normalized gradient describes the local orientation, namely the tangent plane, of the isosurface passing through any point in a volume. The isovalue of this surface is obviously the same as the scalar value at the considered position. Similarly, the curvature of the same isosurface can also be computed and used as an input to a TF.

Different types of curvature measures can be computed at a point on a 2D surface. In the context of volume rendering, the two *principal curvatures* are usually computed. The curvatures consist of the minimum κ_1 and maximum κ_2 principal curvature magnitudes, and if desired, the corresponding principal curvature directions. Based on these measures, other attributes, such as the Gaussian curvature ($\kappa_1 \kappa_2$) or the mean curvature $\frac{1}{2}(\kappa_1 + \kappa_2)$, can be derived as well.

Hladuvka *et al.* [HKG00] have introduced the concept of *curvature-based TFs* with a 2D domain in the context of direct volume rendering, where the curvatures are computed by locally fitting surface descriptions. The curvature TF is a 2D TF that is indexed by (κ_1, κ_2) . More examples of curvature TFs are demonstrated in Kindlmann *et al.* [KWTM03]. Their work presents a significantly simplified computation of the two principal curvature magnitudes κ_1 and κ_2 by using tri-cubic B-spline convolution filters. As Hadwiger *et al.* [HSS*05] show, the isosurface curvature in a volume can also be computed in real-time using tri-cubic B-splines and texture mapping hardware. The authors also compute and visualize the principal curvature directions. The curvature attributes can also be used as additional input to focus and context visualizations, which has been demonstrated by Krüger *et al.* [KSW06] and is further discussed in Section 6.2.1.

4.2. Surface attributes

In most cases, TFs are used to map the scalar or multimodal volumes to optical quantities, and the mapping is used for direct volume rendering. However, several approaches use TFs to analyze surfaces in a volume directly.

To inspect the geometrical variability of isosurfaces in scalar fields with uncertainty information, Pfaffelmoser *et al.* [PRW11] propose to study the mutability in the data using a stochastic uncertainty model. To visualize this model, their isosurface ray-casting approach uses a TF to color the surfaces based on the approximate spatial deviation of possible surface points from the mean surface. Haidacher *et al.* [HBG11] use the TF domain to express surface similarity and dissimilarity for multimodal volumetric datasets by means of information theory measures.

4.3. Statistics and uncertainty

A variety of statistics, usually characterizing local neighborhoods centered at the current voxel of interest, can be very powerful attributes as input to 2D or MD transfer functions. One example is the work of Haidacher *et al.* [HPB*10], who use a statistical domain defined by mean value and standard deviation, in which a 2D TF can be defined. In Section 5, we will describe further approaches that make use of volume statistics in the context of transfer functions.

Kniss *et al.* [KUS*05] describe an approach to interactively explore the uncertainty and risk of surface boundaries. The decoupling of classification and color mapping may point to application specific solutions in situations where many volumes are involved simultaneously. In such cases, the TF interface could remain simple and robust concerning scalability as the feature space is broken down into independent components. Lundström *et al.* [LLPY07] deal with uncertainty visualization through probabilistic animation methods. It could be interesting to explore extensions of the probability representations into the temporal domain and to investigate the general applicability of probabilistic animation. Soundararajan and Schultz [SS15] model the uncertainty in probabilistic TFs resulting from supervised learning approaches. This is an interesting direction to inspect further possible applications of machine learning to volume visualization and TFs, such as adaptive, online, and transfer learning.

4.4. Flow analysis

Even though TFs are commonly used in the analysis of flow data, not much work has dealt explicitly with the design of TFs for flow visualization. Volume rendering is often used as a second step to reduce the clutter in dense texture methods, for instance. An example of this approach is presented by Falk and Weiskopf [FW08], who apply volume rendering to 3D Line Integral Convolution. MD TFs are also of interest in dealing with multiple scalar feature identifiers, and Park *et al.* [PBL*04] suggest using volume rendering with MD TFs. They extract scalar values from the flow field, such as velocity, gradient, curl, helicity, and divergence, and use these values as TF parameters, which results in an expressive and uncluttered visualization. Svakhine *et al.* [SJEG05] present reduced user interaction by allowing only two variables to control color and transparency. The long and cumbersome fine tuning of the transfer function needed in previous work [PBL*04] is thereby avoided. Daniels *et al.* [DANS10] have explored similarities of vector fields by encoding the distances in attribute space. The method is related to explicit TF approaches as it allows users to control paint strokes and color choices on a canvas used for feature enhancement.

A more comprehensive overview of rendering aspects of flow visualization is found in the state-of-the-art report on illustrative flow visualization by Brambilla *et al.* [BCP*12]. The report asserts that volume rendering is known to generate cluttering and occlusion if used unwisely and is not well suited for conveying directional information. They also conclude that volumetric data is often used in flow visualization to show scalar variables such as pressure or temperature.

4.5. Geometric relations

A simple yet highly effective attribute to include in the TF definition is the distance to some geometric entity, such as a point, parametric shape, or arbitrary geometry. Tappenbeck *et al.* [TPD06] propose such an approach and discuss methods for TF specification and suitable applications of it. The efficacy of this approach depends, to a high degree, on the ease of specification, but has the potential to be simplified with the automated generation of a point, or geometry, of reference. A related derived attribute is the feature scale, encoded into a 3D scale field, that expresses the size of the local feature on a per-voxel basis. This approach has been presented by Correa and Ma [CM08]. The method yields convincing images showing distinguished features that are otherwise difficult to classify. Correa and Ma [CM09a] have continued to explore aspects of geometric relations to improve classification of features of interest by introducing the *occlusion spectrum* (Figure 7). The spectrum is based on ideas of ambient occlusion and is able to enhance structures rather than boundaries, as is the case for intensity-gradient magnitude 2D TFs. The ambient occlusion of a voxel is rationalized into a weighted average of the surrounding neighborhood.

Another TF attribute that is useful in some application domains is *orientation* or *direction*. Fritz *et al.* [FHG*09] have computed the orientations of steel fibers embedded in reinforced concrete obtained via industrial CT scanning. The resulting orientations can be parametrized with two Euler angles, where anti-podal points on the unit sphere are identified as describing the same orientation. The

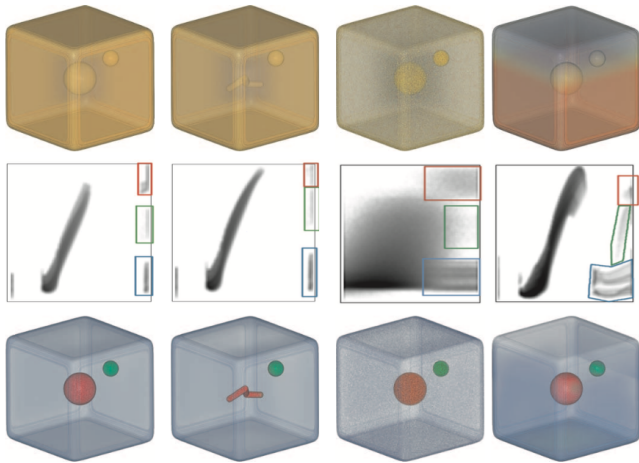


Figure 7: The occlusion spectrum [CM09a] emphasizes the classification of structures within the data, rather than on boundaries between materials. Image courtesy of Correa. Copyright 2009 IEEE.

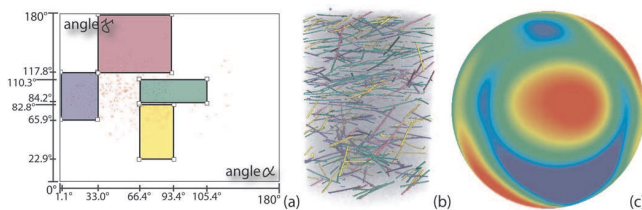


Figure 8: Orientation as an attribute for a 2D TF [FHG*09]: (a) The TF is defined in terms of Euler angles (here, identifying anti-podal points); (b) Specifying widgets in the TF domain selects objects such as steel fibers according to orientation; (c) A directional histogram on the unit half-sphere. Image courtesy of Fritz. Copyright 2009 IEEE.

two angles are then used as inputs to several TF variants, for example a 2D TF including both angles, or a 2D TF for scalar density plus one selected angle, which is important in this particular application. See Figure 8 for an example.

We can view orientation as a specific texture property, and expand the set of attributes to include statistical measures of the textural properties of the data. In Caban and Rheingans [CR08], the classification domain was enriched to support differentiation of similar data values present in different regions based on textural properties in the local neighborhood. Here, the textural properties are defined from various first- and second-order statistics, such as mean, variance, energy, inertia, etc. More advanced textural properties and features of the data, such as shape, as suggested by Pražni *et al.* [PRMH10], require more advanced preprocessing and are therefore considered in Section 5.4.

5. Aggregated Attributes

Conceptually, the introduction of additional quantities in the TF definition makes it possible to discriminate more parts in a dataset.

In practice, however, any added quantity makes it more difficult to create the TF. Visually, 1D TFs can be represented as a line or curve, but in the case of 2D TFs, visualizing and modifying polygons becomes significantly more complex. 3D TFs require volume rendering of the TF itself, which in turn, requires “meta transfer functions” for display, making the approach impractical. To avoid these issues, and going beyond 3D TFs, several papers have been published suggesting approaches for aggregating attributes and reducing the complexity of visualizing and designing higher-dimensional TFs.

5.1. Histogram clustering

To reduce the degrees of freedom in designing a TF, several approaches have been proposed to analyze and cluster the histogram space. User interaction is simplified to select, weight, and modify these clusters. Tzeng and Ma [TM04] propose using the iterative self-organizing data analysis technique to find the clusters in 2D histogram space. With this algorithm, the user can choose opacity and color for each cluster (see Figure 9). Wang *et al.* [WCZ*11] have also worked with 2D histograms but propose modeling the histogram space using a Gaussian mixture model. An elliptical TF is assigned to each Gaussian, and the user interaction is simplified to parameterizing these ellipsoids. Li *et al.* [LZY*07a, LZY*07b] concentrate on industrial CT applications and 1D histograms. They propose using stochastic methods to differentiate between the clusters in the histogram.

Maciejewski *et al.* [MWCE09] describe a method for 2D clustering in a feature space, comprised of value and gradient. The technique automatically generates a set of TF components that can be refined or filtered out later in the workflow. The method also incorporates temporal changes by building a histogram volume from the 2D feature space, which allows for a more consistent clustering over time.

In a more general approach, more suitable for arbitrary attribute spaces, Wang *et al.* [WZK12] use hierarchical clustering. From a preprocessing step, the user can select clusters in terms of subtrees in the modified dendrogram, and subsequently refine the TF selections in a more fine grained multidimensional space. Recent work by Ip *et al.* [IVJ12] demonstrates an algorithm for multilevel segmentation of the intensity gradient 2D histogram that allows the user to select the most appropriate segments hierarchically for which they can generate separate TFs.

5.2. Local frequency distribution

Lundström *et al.* [LLY05, LLY06a] propose a set of techniques and tools based on local frequency distributions (LFDs). An exhaustive peak finding approach aims to find all neighborhoods, data blocks, to associate with a peak in the global histogram. To this end, the partial range histogram (PRH) is introduced in which each neighborhood has its LFD footprint measured against a range that is automatically generated from the current global peak searched. Neighborhoods with a sufficiently large LFD footprint in this range are added to the group, making up the PRH. After each iteration, the LFDs of the current PRH are removed from the global histogram, revealing another peak, and the process is iterated until all neighborhoods have

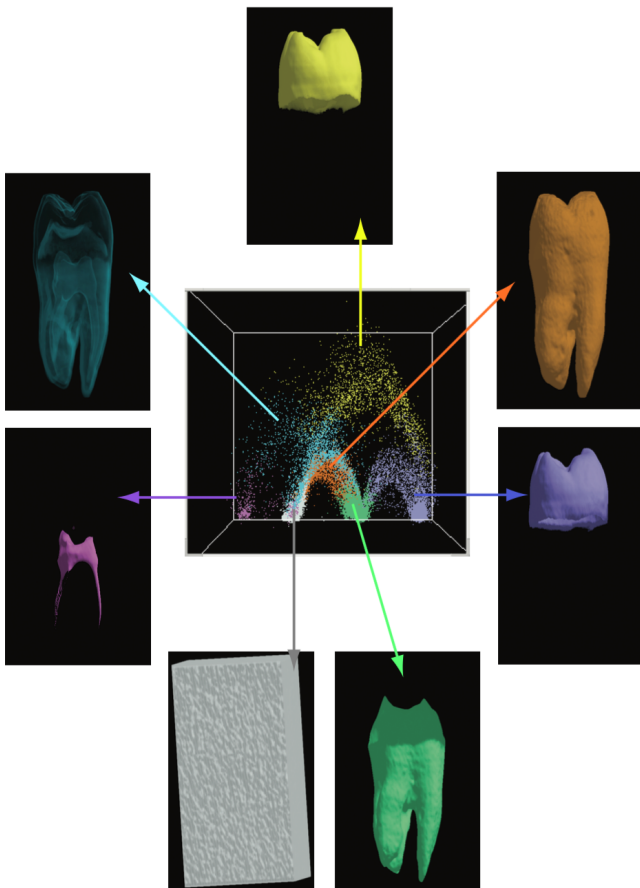


Figure 9: The clustered histogram is shown in the center. The volume rendered images depict the corresponding regions in the spatial domain [TM04]. Image courtesy of Tzeng and Ma.

been assigned. In the second part of the process, tissue classification is aided by the neighborhoods' LFD footprints, which are used in a competitive classification approach, to determine probability and tissue classification. An example of this approach is shown in Figure 10. In Lundström *et al.* [LYL*06], LFDs are enhanced to appear more clearly in the global histogram, which improves peak visibility in the user interface. This enhancement is achieved by promotion of local correlation by simply applying a power exponent, α , to the binned values in the local histogram. Promoting local peaks is also helpful in peak detection and in automated TF adaption to new datasets. The method was evaluated by experts on both CT and MR data. The α -histogram is not as successful as the PRH approach in adapting TFs.

Serlie *et al.* [STF*03] have introduced the concept of *LH histograms* for identifying material boundaries in the context of virtual colon cleansing in virtual colonoscopy. For each voxel, a respective *low* (L) material value and a *high* (H) material value are determined in the positive gradient direction for H, and in the negative gradient direction for L. The voxel then corresponds to a point in the 2D LH space, which denotes a material transition between a material with value L and a material with value H, respectively. Šereda *et*

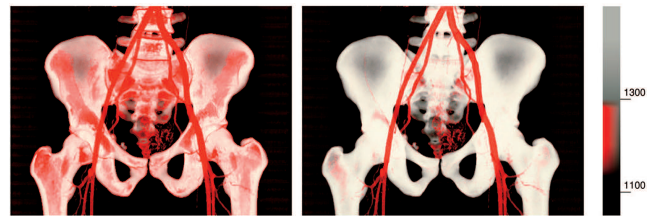


Figure 10: Separation of spongy bone and vessels. The left image shows the results with a standard 1D TF, which renders bone and vessels red, but a classifying 2D TF makes the vessels stand out from the background [LLY06a]. Image courtesy of Lundström. Copyright 2006 IEEE.

al. [SBSG06] have extended this concept to general transfer functions whose domain is the 2D LH space. They also introduce the use of hierarchical clustering in the LH space for semi-automatic TF design [ŠVG06].

To better capture the neighborhood properties around voxels, Patel *et al.* [PHBG09] define the concept of *moment curves*. These curves describe the moments in the first- and second-order statistical attributes for each voxel. These attributes depict a point projected back into a 2D space where TF primitives are defined, which yields a more powerful way to identify features of interest.

Johnson and Huang [JH09] have implemented the concept of local distributions in a very general and query-driven approach, in which they define *bins* and *clauses*. With a bin, the user restricts the domain of the distributions' given intervals. Clauses are series of Boolean predicates that compare bin variables to other quantities. Both bins and clauses are defined in a domain-specific language and can be changed arbitrarily. The color for each voxel is determined by its predicate matchings and the opacity by the quality of this matching.

Lindholm *et al.* [LLL*10] have used LFDs to support a logic framework based on labeling the LFDs as blood, gas, iodine, liver, etc., which allows the creation of expressions to render iodine when close to, or not close to, liver tissue, for instance. These labels provide for the semantics of the TF components, and the logic is then executed on the GPU during rendering.

Similar to the approaches of Johnson and Huang [JH09] and Lindholm *et al.* [LLL*10], where the TF specification is enhanced by rules, Cai *et al.* [CNC015] propose a rule-enhanced TF as well. In this approach, the rules have been determined through training on segmented datasets to identify different tissues in the data when derived data attributes and LFDs are used.

5.3. Dimensionality reduction

Dimensionality reduction methods propose simplifying the complexity of the TF design process by projecting high-dimensional TFs to a lower-dimensional space. Kniss *et al.* [KUS*05] have presented an approach for statistically quantitative volume visualization in which a graph-based technique is used to achieve dimensionality reduction. In this work, uncertainty visualization and probabilistic classification take a central role (also see Section 4.3.)

De Moura Pinto and Freitas [dMPF07] propose using unsupervised learning, specifically in the form of self-organizing Kohonen Maps to detect structure in the data. These maps are trained with input voxel values and local derived quantities (such as derivatives and statistical measures). Linear combinations of cells in the Kohonen Map are then assigned to the corresponding voxels in the volume dataset, and the TF is applied to that map instead of the high-dimensional attribute space.

Haidacher *et al.* [HBKG08] aim to reduce multivariate data into a single fused representation, which is then mapped via the well known value/magnitude 2D TF. The reduction is carried out using weighting based on point-wise mutual information. Kim *et al.* [KSC*10] follow a similar approach but use Isomap, local linear embedding, and principal component analysis to reduce the high-dimensional MD-TF/multichannel data space and allow for straight 2D TF application. Zhao and Kaufmann [ZK10] combine the local linear embedding method to reduce the dimensionality with a user interface based on parallel coordinates. Dimensionality reduction using Isomap has also been applied in Abbasloo *et al.* [AWHS16] for the visualization of tensor normal distributions.

5.4. Topology and skeletons

The use of topological methods has recently been effective in visualization and analysis of many types of data. These methods provide aggregated attributes that can be useful for TFs. In their early works, Fujishiro *et al.* [FAT99, FTAT00] used a hyper Reeb graph to automatically generate TFs that emphasize critical isosurfaces, which are surfaces corresponding to isovalues near a change in the surface topology. Takahashi *et al.* [TTF04] improve this approach by using topological volume skeletonization, which improves the detection of critical field values and speeds up computation in comparison with hyper Reeb graphs. The volume skeleton tree concept is also used by Takeshima *et al.* [TTFN05] to generate transfer functions using inclusion levels, isosurface trajectory distances, and isosurface genera. This concept is further extended by Weber *et al.* [WDC*07] who give the user more control over the TF generation by enabling the specification of different TF parameters per topologically distinct feature. Furthermore, Zhou and Takatsuka [ZT09] use contour trees to partition the volume in subregions. These authors take particular care to automatically assign matching colors and opacities to those subregions. Instead of considering per-voxel features, Praßni *et al.* [PRMH10] have suggested assigning shape properties to features, which aids the user in assigning material properties for different volumetric structures. The different structures are identified using curve-skeleton analysis in a preprocessing step. Lastly, Xiang *et al.* [XTY*11] have introduced the skeleton-based graph cuts algorithm that enables effective and efficient classification of topological structures used to assign localized transfer functions.

5.5. Domain specific aggregation

Although many papers have been written with a specific application domain in mind, such as medicine or engineering, most of the presented techniques are universally applicable. In some cases, however, the data exhibits very subtle and domain-specific structures that require consideration through highly specialized methods.

Diffusion tensor imaging (DTI) is one such specific domain, and DTI data is often visualized using other visualization techniques, such as glyph-based or with feature extraction approaches that include DTI fibers. Kindlmann and Weinstein [KW99] present a DVR method in which the tensor field is colored using *hue-balls* and shaded with *lit-tensors*. A 2D barycentric space of anisotropy is used to define the opacity of samples. In Kindlmann *et al.* [KWH00], the approach is studied in more detail, and the options are extended to include barycentric color maps.

In more recent work, Bista *et al.* [BZGV14] present an approach for volume rendering of diffusion kurtosis imaging (DKI) data that more clearly depicts microstructural characteristics of neural tissues. Spherical harmonics are used to color and shade samples in the volume rendering based on the spatio-angular field of DKI.

Seismic data is another domain with specific attributes that constitutes many dimensions and thus requires advanced user interfaces to control. Zhou and Hansen [ZH13, ZH14] present examples of this, which are discussed in Section 8.4.1.

Another approach, which resembles textural classification methods, or texture-based TFs, has been proposed by Alper Selver [Sel15]. Brushlet expansion is applied to the original volume and allows the analysis of the resulting quadrants in order to identify low- and high-frequency textures. By selecting quadrants and thresholding of the complex brushlet coefficients, specific textural properties can be reconstructed in the volume. The quadrant selection can be manual, atlas based, or based on machine learning.

6. Rendering Aspects of the Transfer Function

So far we have reviewed the dimensionality aspects of TFs, including discussions on global and local attributes (direct, derived, and aggregated) that aim to improve the classification power of the TFs and to determine colors and opacities. In this section, we review publications that deal with applying the TF and the rendering of the volume. Some of the techniques discussed in these publications deal with unique data, such as the full-color Visible Human, or use TFs to achieve specific artistic or stylistic effects. Other publications deal with enabling focus and context visualization by using specific methods and TF concepts.

6.1. Coloring and texturing

One interesting way of computing a TF is by attempting to apply realistic looking colors to gray level volume data by either coloring voxels or utilizing texture patterns. Realistic looking volume renderings can be obtained by *transferring* the colors of a colored volume such as the photographic Visible Human dataset (obtained from actual photographs) to a volume of a different modality. In this way, realistically colored volumes can be rendered, although no actual color information is known. A powerful approach for transferring colors is to train a neural network for this purpose, one example being the work by Muraki *et al.* [MNKT01]. The authors transfer colors from the Female Visible Human dataset to MRI data by training an RBF network.

Instead of transferring colors on only a per-voxel basis, larger *texture patterns* can also be transferred. Approaches in this direction

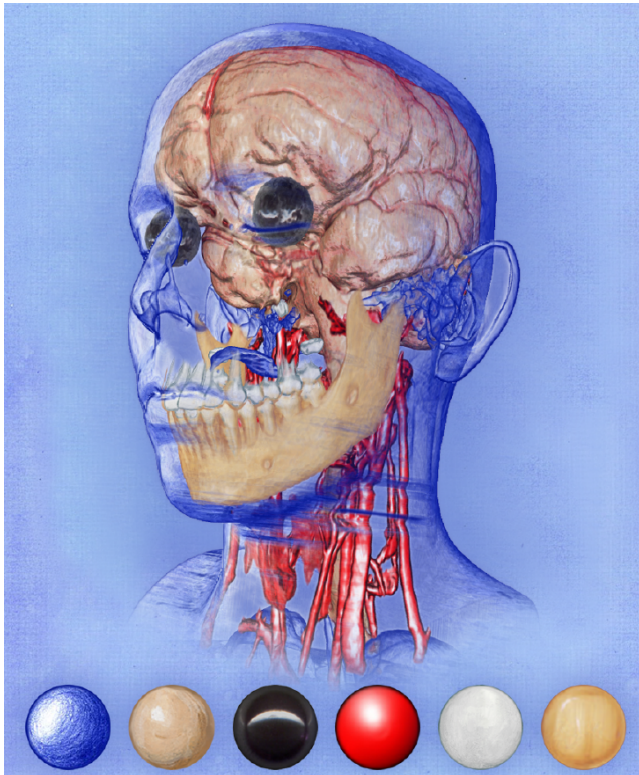


Figure 11: Segmented volume rendered with a MD style TF based on data value and object membership. The base of the image depicts the lit spheres for the different styles. Image courtesy of Bruckner et al. [BG07].

differ by transferring patterns from either 2D texture maps or 3D texture maps. For virtual colonoscopy applications, Shibolet and Cohen-Or [SCO98] present a two-step approach to map patterns from 2D textures to the 3D non-convex surface of the colon. In this way, a more realistic look of the colon is obtained. Lu and Ebert [LE05] describe a system for obtaining example-based volume illustrations. They use colored example images to compute a colored 3D volume with a similar look using texture synthesis. They then transfer this look to a volume for rendering either slices or the whole volume.

Traditional color and opacity TFs can be combined with the application of 3D texture patterns to the volume, as has been performed by Manke and Wünsche [MW09] using texture TFs. This combination results in what they call *texture-enhanced direct volume rendering*, which can be used to visualize supplementary data such as material properties and additional data fields.

6.2. Rendering styles and shading

The general idea of a TF can also be extended to determine the *rendering style* according to volume properties. Approaches range from determining or modifying the shading model that is used, to general illustrative, non-photorealistic volume rendering techniques.

A general concept is that of *style transfer functions*, which have been introduced by Bruckner and Gröller [BG07]. Similarly to a traditional TF, optical properties are assigned according to voxel density. However, the optical properties are not simple colors and opacities, but instead entire rendering styles. The different styles are described using an image-based lighting model built on lit sphere maps (Figure 11). Different styles can be interpolated, which allows for combining them into a single illustrative volume rendering.

This concept was taken further by Rautek et al. [RBG07], who have introduced the concept of determining a TF using *semantic layers* by mapping several volumetric attributes to multiple illustrative visual styles. Semantic layers allow a domain expert to specify this mapping in the natural language of the domain using a linguistic specification. This specification is then mapped to the visual style used for volume rendering by employing fuzzy logic techniques. This concept was later extended to include interaction dependent semantics [RBG08].

Csébfalvi et al. [CMH*01] have proposed a model to render contours of volumetric objects by combining two measures in a way that is similar to a gradient magnitude-weighted TF. The first measure ascertains how much a voxel corresponds to a surface, which is determined from the gradient magnitude. The second measure establishes how much a voxel corresponds to the silhouette of an object, which is determined from the angle (via the dot product) between the gradient direction and the view direction. Both measures are then multiplied to determine contours within the volume without requiring an explicit correspondence to isosurfaces.

Maximum intensity projection (MIP) is an alternative to compositing colors and opacities from a TF that uses the emission-absorption volume rendering integral given in Section 2. MIP retains only the maximum value encountered along each ray. This approach provides the advantage of avoiding the need to specify a full TF. However, MIP has the disadvantage of losing occlusion and depth cues. A way to circumvent this disadvantage and to combine the advantages of MIP and classical DVR approaches has been proposed by Bruckner and Gröller [BG09] with the maximum intensity difference accumulation (MIDA) approach. By using MIDA, it is also possible to obtain a smooth seamless transition between MIP and DVR.

As presented by Hernell et al. [HLY07], the TF can also define multiple rendering properties. Here not only the traditional absorption/scattering property is used but also light emission, where selected materials are treated as light sources as well. This method works in combination with the ambient occlusion and global illumination approaches.

6.2.1. Focus and context techniques

TFs that contain information about segmented objects can be efficiently employed for *focus+context* approaches. The user's attention should be directed toward the focus, but the context should still be available in a less emphasized and unobtrusive way.

As introduced by Viola et al. [VKG04, VKG05], by assigning different parts of a volume, a corresponding object importance value leads to the concept of *importance-driven volume rendering*. These importance values are then used in the volume rendering to establish a view dependent priority regarding the visibility of different

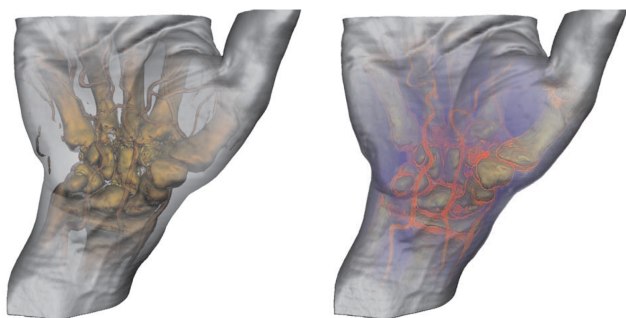


Figure 12: Different combinations of context layers for focus+context visualization in ClearView [KSW06]. Image courtesy of Krüger. Copyright 2006 IEEE.

objects. Doing so allows for an easy distinction to be made between objects in the focus and objects in the context, corresponding to more important and less important objects, respectively.

Another approach involving focus and context is outlined in the work by Hadwiger *et al.* [HBH03] who present GPU-based two-level volume rendering of segmented volume data. In this work, different segmented objects can have different TFs, different rendering modes (such as DVR or MIP), and different compositing modes. The latter capability enables *two-level volume rendering*, which comprises one local compositing mode per object, and a second global compositing level that combines the contributions of different objects. This scheme can be used to distinguish between focus and context, for example rendering the context using the volumetric object-contour technique proposed by Csebfalvi *et al.* [CMH*01] and rendering the focus with standard DVR and a regular 1D TF.

Illustrative approaches can be integrated with focus and context. Bruckner *et al.* [BGKG05, BGKG06] employ a modification of the compositing equation to distinguish focus from context for illustrative context-preserving volume rendering. This strategy enables the efficient simultaneous visualization of interior as well as exterior structures in a volume dataset. The *VolumeShop* system developed by Bruckner and Gröller [BG05] combines many different focus+context and illustrative shading techniques into an application for general interactive illustrative volume rendering. Different segmented objects can have different TFs and different shading modes. Cutaway views and ghosting are other powerful methods for obtaining volumetric illustrations.

Another approach for focus+context rendering is blending multiple focus and context layers into a single output image, as carried out in the *ClearView* system of Krüger *et al.* [KSW06]. See Figure 12 for examples of combining different layers.

6.2.2. Domain-specific examples of stylistic approaches

In some domains, a specific visualization style has been developed. Domain scientists have developed excellent skills in interpreting these types of illustrations over multiple centuries.

Medicine is one well known domain-specific example, and not unexpectedly, a large number of papers cited in this STAR use medical data as examples. However, other research areas also require

specialized methods, such as in the work of Fritz *et al.* [FHG*09] concerning the inspection of steel fiber reinforced sprayed concrete. Another popular field for illustrative volume rendering is earth science, in particular the visualization of seismic data. Patel *et al.* [PGTG07, PGT*08] have presented an algorithm to generate textbook like seismic illustrations. A number of layers are combined, such as 2D textures and 3D volumes, and the 3D volumes are blended with the 2D textures by 1D TFs.

6.3. Animation and temporal techniques

Animations serves as a natural means to convey time-dependent data. Moreover, adding animations can be a very powerful technique to provide insight into volumetric data, even if the data are not time dependent.

One way of using animation techniques is to expose the uncertainty in the data. Lundström *et al.* [LLPY07] have tackled this important problem by mapping the data uncertainty to an interactive animation of the volume. The TF is decomposed into material components. Each component can be given a probability separate from the material property opacity. Material mappings are then animated to convey the probability of the material, that is, the probable extent of the material.

Another approach to incorporate animations is to morph between TFs, a technique proposed by Wong *et al.* [WWT09]. The user specifies a start-TF and an end-TF and the system automatically interpolates all intermediate TFs to obtain a full animation.

The reverse concept, taking temporal data and mapping it into a static view, is presented by Balabanian *et al.* [BVMG08], who introduce temporal styles, so called style TFs, for time varying volume data. Their method condenses multiple time steps of a time varying dataset into a single view, in which the TF provides different styles at different time steps, allowing for the depiction of start/end conditions and internal transition points.

6.4. Sample reconstruction

The way in which a TF can or should be applied to volumetric data often depends on sampling and function reconstruction considerations. Two important considerations are: 1) how should a TF be applied to segmented data when considering the *boundary* between different materials or segmented objects; and 2) how should a TF be applied to down sampled volume data in the context of *multiresolution volume rendering*.

A feature aware approach for applying multiple 1D TFs in an accurate way to a volume containing multiple materials is presented by Lindholm *et al.* [LJHY14]. This approach achieves a better preservation of feature boundaries than previous work. The approach estimates the local support of materials before performing multiple material specific reconstructions. This estimation prevents much of the misclassification traditionally associated with transitional regions when TFs are applied. A result example is shown in Figure 13.

Younesy *et al.* [YMC06] describe a subtle problem in multiresolution volume rendering: the TF should be applied to the *distribution* or *histogram* conceptually associated with each voxel in a down-sampled volume, instead of to the down-sampled voxel value. In

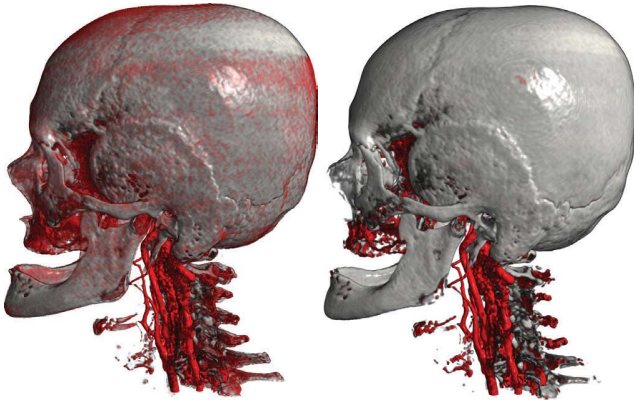


Figure 13: The left image shows the result of standard volume rendering with continuous reconstruction. On the right, the red “hull” is avoided by continuous reconstruction within each feature while preventing interpolation between features as proposed by Lindholm et al. [LJHY14]. Image courtesy of Lindholm. Copyright 2014 IEEE.

order to restrict memory usage, these authors substitute the histogram associated with each voxel with the corresponding mean and variance. The TF is then applied to pairs of mean and variance for each voxel, which leads to an improved quality in multiresolution volume rendering.

The same problem is addressed by Sicut et al. [SKMH14], who apply the TF to a whole distribution of voxel values. However, instead of being restricted to the mean and variance of a distribution, their goal is to represent the distribution accurately, while simultaneously keeping memory usage low. This aim is achieved by employing a sparse representation with 4D Gaussian basis functions in a combined 4D space composed of a 3D volume domain plus a 1D TF domain.

In addition to representing data distributions in multiresolution volumes, a histogram for each voxel can also be used to encode uncertainty, such as the concept of *hixels* that is presented by Thompson et al. [TLB*11]. Hixels provide a histogram of data values for each voxel, which can be employed in the analysis and rendering of large-scale volume data.

Another way to take on the problem of multiresolution volume rendering is from the perspective of defining which optical properties should be down-sampled in order to obtain lower resolutions while retaining good quality. Kraus and Bürger [KB08] have compared down-sampling RGBA volumes, applying an RGBA TF, to down-sampling the corresponding extinction coefficients instead, which led to better results.

6.5. High dynamic range rendering

In a fashion similar to rendering in general, volume rendering can be extended from a low dynamic range to a high dynamic range (HDR) rendering. The work of Yuan et al. [YNPC05, YNPC06] extends this concept to specifying opacities in the TF with higher *precision*, in addition to using a larger range for the colors. As outlined in

Kroes et al. [KPB12], a more general volume rendering pipeline has been developed to support full Monte Carlo path tracing, which also works with HDR.

7. Automation of Transfer Function Definition

Consulting the volume rendering and TF-related literature of the last decade reveals a motivational statement that can be found in many of the abstracts: *Volume rendering has established itself as a powerful visualization tool in many domains, but the necessary design of TFs is a time consuming and tedious task.* Consequently, a significant amount of work is dedicated to the automatic and semi-automatic generation of TFs. Based on the fundamental notions and classifications of the TF concept, this section provides an overview of the higher level aspects of TFs. This section also examines the improvement of functionality and efficiency by introducing automation through data driven approaches as well as by utilizing user knowledge encoded through interaction.

7.1. Adapting presets

Instead of starting the TF design process anew for every new dataset, Rezk-Salama et al. [RSHSG00] have proposed the reuse of existing, manually designed TFs from previous similar datasets and adjusting them by non-linear distortion to fit the new data. To perform the distortion, the work has compared both the histograms of the datasets and improved the initial results using Kindlmann and Durkin’s position function [KD98]. Castro et al. [CKLG98] have suggested to work with labeled components to further facilitate adjusting of presets at this component level. Rezk-Salama et al. [RSKK06] later refined their approach to utilize existing TFs. They have recommended establishing a domain-specific semantic model of components in the data and the manual definition of one or more 2D TF primitives for each component. For a number of datasets, these primitives are adjusted manually to account for the variation between datasets. The adjustments are projected to one dimension using the first principal component, and this single parameter is exposed to the novice user as a slider.

In extension of their previous work on LiveSync, a system to synchronize 2D slice views and volumetric views of medical datasets, Kohlmann et al. [KBKG08] have presented a new workflow and interaction technique to aid the TF generation. As a starting point, the user selects a point of interest (POI) in the slice view, and the system refines and adapts TF presets based on statistical properties derived from the selected POI. The system provides additional supportive features, such as smart clipping, to reveal volume views that are not occluded.

7.2. Semi-automatic generation

Pursuing automatic and semi-automatic generation of TFs is the ultimate goal in many application domains since they enable a more widespread use of volume rendering. In most situations, it makes sense to allow the user to change the TF manually even if it was *generated* automatically. Consequently, we distinguish between automatic and semi-automatic methods by their capability to generate an (initial) TF with or without user intervention.

Many different approaches have been used over the years to accomplish the goal of a semi-automatic TF generation. He *et al.* [HHKP96] were the first to propose a semi-automatic approach to TF design in 1996. They employed a “supervised genetic algorithm” approach to find the best TF with either the user or the system itself, automatically ranking generations to identify the fittest members for the next iteration. Early work on semi-automatic generation also includes the work by Castro *et al.* [CKLG98] who have proposed a component-based weighted mixture to yield the complete TF. Components can then be referred to at a higher abstraction level as different tissues, thus providing a semantic description. In adopting a novel view to the TF as a two-step process, Fang *et al.* [FBT98] have described a semi-automatic algorithm. In their implementation, they first apply a transformation to the entire dataset using image enhancement and boundary detection operations, followed by a simple 1D linear ramp to “color” the dataset.

Leveraging histograms is a common method for the semi-automatic generation of TFs. While focusing on material boundaries in datasets, Kindlmann and Durkin [KD98] have considered the 3D histogram volume defined by attribute value, first and second directional derivative. In this space, they have defined a model that has classified boundaries and then have used this model for automatic opacity function generation. Roettger *et al.* [RBS05] have used a 2D value/gradient histogram and x, y, z spatial coordinates of the corresponding values to classify the dataset. Prauchner *et al.* [PFC05] have combined the Design Galleries [MAB*97] concept with the histogram approach of Kindlmann and Durkin to evaluate detailed views of the TF, thus guiding the user in finding suitable TF settings. Correa and Ma [CM09b] have aimed to automatically optimize the visibility of certain value ranges. The ranges are selected by the user through modulating the opacity to ensure that otherwise hidden features are visible. The approach uses a visibility histogram along the rays. Correa and Ma have later extended their work [CM11] by adding another dimension to the visibility histogram and by also extending the algorithm to a view independent iterative mode.

Clustering with feature space is another method for the semi-automatic generation of TFs. Šereda *et al.* [ŠVG06] have employed an approach based on agglomerative hierarchical clustering in the two-dimensional LH domain [ŠBSG06] for identifying material boundaries. The clustering can be computed according to two alternative similarity measures. The first measure clusters similar boundaries, which is determined by their distance in LH space. The second measure clusters boundaries that are spatially connected in the volume. The resulting cluster hierarchy can be explored by the user to facilitate semi-automatic TF design. Maciejewski *et al.* [MWCE09] have proposed a method for 2D clustering of the feature space (value, gradient) that automatically generates a set of TF components (see Section 5.1). Specifically for abdominal visualization, Selver and Guzelis [SG09] have designed a method that creates a histogram for each slice of a scan and puts those together to create a volume histogram stack. The peaks are then extracted automatically via an RBF network. Nodes are hierarchically generated at different levels, which can be selected to form groups to which the user assigns TF materials.

Rather than generating the TF semi-automatically, some authors propose a semi-automatic transfer of properties. Pražni *et*

al. [PRMH10] have related shape properties to features, rather than to voxels. This approach aids the user in the assignment of material properties for different volumetric structures. The different structures are identified using a curve-skeleton analysis in a pre-processing step. Ip *et al.* [IVJ12] have demonstrated an algorithm to segment the 2D histogram and allow the user to select the most appropriate segments hierarchically for which they can generate separate TFs.

With MD data, data with multiple attributes at each voxel, projecting the high-dimensional space to lower dimensional spaces formed by combinations of attributes can guide the user in the selection of an appropriate TF. Such multivariate data has been addressed by Liu *et al.* [LWT*14] who have used dynamic and animated projections to guide the user through the attribute space. The user is presented with multiple projected views in the TF design process.

7.3. Automatic generation

The automatic generation of TFs does not require user interaction. For the visualization of uncertainty of isosurfaces, Pfaffelmoser *et al.* [PRW11] have used an automatically generated TF to color the surfaces based on the approximate spatial deviation of possible surface points from the mean surface (see Section 4.2). The idea of using Morse theory to automatically decompose the feature space of the volume into a set of cells has been presented by Wang *et al.* [WZL*12]. Using cell separability, cells that potentially contain important features are retained, whereas those most likely originating from noise are rejected. The remaining cells are transformed into a hierarchical representation by successively merging them. Finally, the user is presented with these cells and can assign color and opacity values to the individual cells or groups of cells. Fujishiro *et al.* [FAT99] and Fujishiro *et al.* [FTAT00] have automated the color assignments for the TF with a starting point in 3D field topology and graph representations. Wang and Kaufman [WK12] have used an importance function that the user can modify to automatically compute the TF. The importance of objects in the volume is conveyed using the saliency of a color. Bramon *et al.* [BRB*13] have used information theoretic strategies that are based on concepts, such as informativeness and informational divergence, to automatically define multimodal TFs.

7.4. Supervised machine learning

The design of a “good” TF requires some experience and domain knowledge. An interesting research question that arises is if this knowledge can be gained through machine learning. Tzeng *et al.* [TLM05] have proposed to do exactly that. They have suggested to let the domain scientist interact with the materials in a volume directly and hide the underlying classification through the use of machine learning. In this work, the user paints onto certain areas (similar to a segmentation user interface), and the system uses machine learning to determine a multidimensional TF based on the input. They claim that almost any machine learning technique, such as artificial neural networks, support vector machines, Bayesian networks, or hidden Markov models, is applicable. The authors implemented neural networks and support vector machines. Both showed promising results without one being superior over the other. Inspired

by Tzeng *et al.* [TLM05], Soundararajan and Schultz [SS15] have provided a more recent overview in which they studied the use of a number of different classifiers. The conclusion from their work is that the random forest approach is the recommended classification technique. De Moura Pinto and Freitas [dMPF07] instead have proposed the use of unsupervised learning methods (Kohonen Maps) to automatically reduce the dimensionality of the TF space and assign color and opacity to the resulting low-dimensional space. Wang *et al.* [WCLC06] have assumed a set of parameters from expert input that classifies volumes, termed an *information bank*. For a new dataset, they used Kohonen Maps to find similarities with the information bank. They generated a TF using the known entries from the information bank and adjusted those to the current dataset using a back-propagation network.

8. User Interfaces

Interactive volume rendering is a powerful tool for visual exploration of volumetric data. Interaction and exploration are thus key aspects that converge in the user interface for the TF. As stated in the introduction, balancing between simplicity and classification power has been the driving force in a substantial body of research in this field. In this section, we discuss the user interface aspects of TFs and provide an overview of what has been accomplished so far.

8.1. One-dimensional editors

For the direct volume rendering task, the colors and material properties are commonly compiled into textures and stored on the GPU where they can be efficiently applied in the process of making an image. In a simple approach, the user can place control points, or draw a curve, to define colors and opacities, an example of which is shown in Figure 3. At each control point, a color is assigned, and the vertical position is used to set the opacity. An early example of such an editor is found in Cignoni *et al.* [CMPS97]. A year later, Castro *et al.* [CKLG98] have published work on the specification of TFs for medical data by combining separate components, each component representing a tissue of interest. A component was expressed as an envelope, or a primitive such as a trapezoid, and the combination of all components resulted in the final lookup table. Castro *et al.* carefully studied the effect on the final result for a range of parameters defining the exact shape of the component. König and Gröller [KG01] also developed a component-based editor for the VolumePro hardware.

Histograms are useful for the aggregation of attributes as described in Section 5. They can also aid the user. For example, the global frequency distribution is often shown as a histogram in the background of the editing viewport. The global histogram is sometimes of moderate use and provides poor guidance as the features of interest are often non-visible peaks, drowning in the overall mix of materials. One idea proposed by Lundström *et al.* [LYL*06] is the α -histogram, which is able to emphasize locally homogeneous regions in the global histogram (see Figure 14). The Visibility Histogram, introduced by Correa and Ma [CM11], is another approach to help the user define a TF by providing feedback of what information is lost in the volume for the current TF. Potts and Möller [PM04] also studied 1D TF editing. They suggested a logarithmic scale on

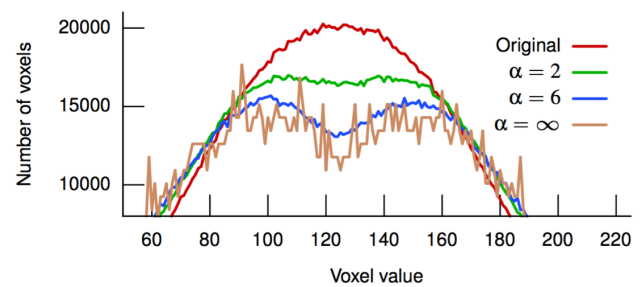


Figure 14: Two Gaussian distributions that combined completely hide the separate peaks in a normal histogram whereas the α -histogram is able to recover the peaks [LYL*06]. Image courtesy of Lundström.

the opacity axis, which better aligns with visual appearance in the rendered image. Maciejewski *et al.* [MJW*13] proposed the abstract attribute space, a histogram showing a selected statistical measure such as mean, standard deviation, or skewness. This space directs the user towards peaks in the histogram with greater interest, compared to the standard frequency histogram. The authors also show how the abstract attribute space can be helpful in 2D density plots.

Often times for simulation data, floating point values are provided. This defines a much larger range than typical color and opacity ranges, which use 8-bits per channel. When the data range is greater than the resolution of the color channels, as in HDR rendering (see Section 6.5), it is necessary to provide a user interface to allow for the specification of opacities and colors in the higher range of the data. A user interface to control this mapping can aid the user in highlighting details in the data. Yuan *et al.* [YNCP05, YNCP06] have proposed to use nonlinear scaling and zooming of the large HDR TF domain to allow users to adjust the TF even for small scale features.

8.2. 2D and multidimensional editors

Taking the step from one-dimensional TF editors to two-dimensional and multidimensional TF editors undoubtedly raises the complexity drastically. Many users find that interacting with 1D TFs is difficult, yet the classification power is not expressive enough for many domain needs. In this subsection, we review the literature on 2D and MD TFs and associated user interfaces.

Non-separable 2D TFs have mostly been supported by an editor exploiting a 2D space. Analogous to the 1D TF editor, and even more so, the two-dimensional frequency distribution of data values becomes an essential aid in the user interface. As the two dimensions of screen space are taken by the 2D data value domain, the 2D frequency distributions of the data domain need to be represented as 2D histograms, or density plots, rather than 1D traditional histograms. Thus, in the discussion of 2D TF editors, both TF components and 2D histogram are tightly coupled. Kniss *et al.* [KKH02] have presented a set of widgets and interactions tools to define the 2D TF. The gradient magnitude is most commonly used for the second dimension, and the widgets presented are mostly geared towards that type of data (see Figure 15). A set of 2D widgets, or parametrized

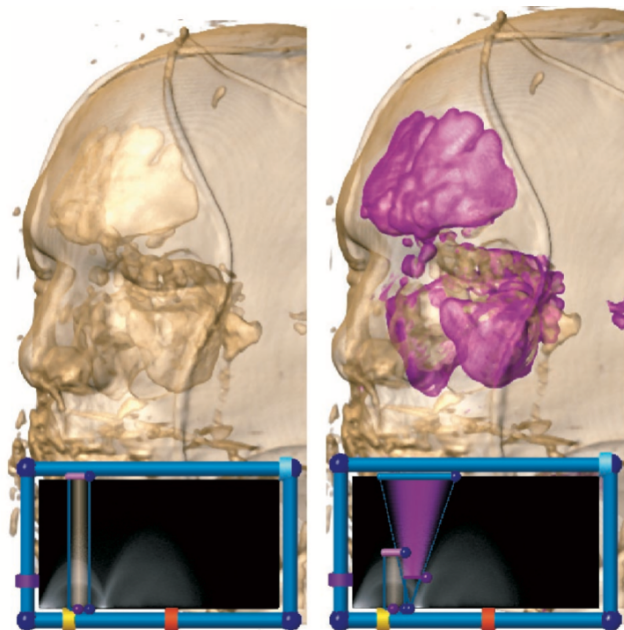


Figure 15: The two-dimensional TF editor widgets allow the user to more precisely control the classification power of the 2D TF. The triangular widget can also be skewed to better fit the desired classification domain [KKH02]. Image courtesy of Kniss. Copyright 2002 IEEE.

primitives, for 2D TF editors has also been presented by Rezk-Salama *et al.* [RSKK06] Zhou *et al.* [ZSH12] presented a technique for combining 2D and 1D TF editors, in a side-by-side manner, with the goal to increase the expressiveness while minimizing user interface complexity. They argued that users have familiarity with 2D and 1D TF editors, and combining them to achieve 3D TFs provides a more effective solution without overwhelming users. A user is given the choice to select different 2D domains and an optional 1D TF domain where a TF curve is defined for each widget created in the 2D domain. The user can specify selections using widgets in the form of rectangles or freehand drawn areas using a lasso tool.

Another approach for user interfaces for high-dimensional TFs, Lundström *et al.* [LLY06b] introduced the *sorted histograms*, for 2D and MD editors. Their technique maintains the simpler visual approach along the style of the 1D TF histogram representation. The additional attribute is color-coded and stacked in a form similar to that of the 1D histogram. Selections on the value range can be done iteratively to cover additional dimensions of the TF.

8.3. Parallel coordinates

Reaching beyond 2D TF editing becomes highly intricate for any direct primitive-based editing, such as an envelope or areal shape. Several researchers have proposed Parallel Coordinates (PC) to deal with TF specification. In Akibay and Ma [AM07], this visualization technique has been used in combination with temporal density plots and 2D histograms, to show the temporal development for time varying multivariate data. Their user interface has supported

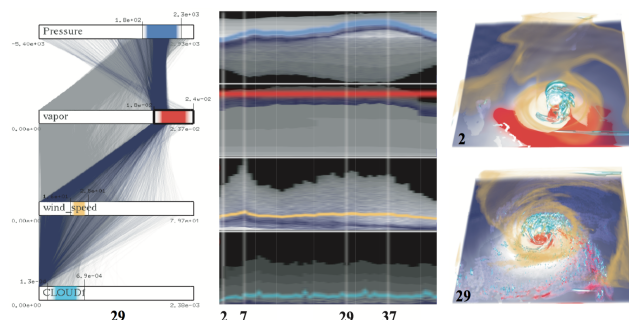


Figure 16: Use of parallel coordinates to make range selections of TF components for multivariate volumes. The center view shows a temporal 2D histogram [AM07]. Image courtesy of Akibay.

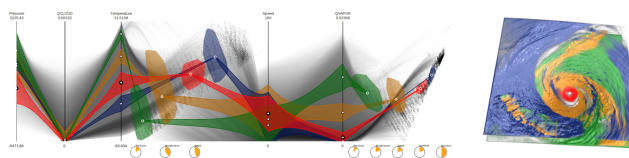


Figure 17: Parallel coordinates with embedded multidimensional scale plots for editing multidimensional TFs [GXY11]. Image courtesy of Guo. Copyright 2011 IEEE.

simultaneous views of two time steps and brushing over the range in the PC view (see Figure 16). In Zhao and Kaufman [ZK10], the parallel coordinates have been used in addition to a dimension reduction approach that employs a local linear embedding to simplify the TF design. Dimension reduction has also been used by Guo *et al.* [GXY11] in combination with PC plots. They have integrated multidimensional scale plots within the parallel coordinates plot to allow the user to edit the TF settings in a single consistent view. An example is shown in Figure 17. Guo *et al.* [GXY12] later extend their previous work to combine PC plots with multidimensional scaling to avoid the need to adjust parameters for each dimension separately. Liu and Shen [LS16] have proposed to use association analysis for multivariate scientific data sets to generate a parallel coordinate plot. They introduce the concepts of uniqueness and informativeness to select scalar values of interest.

8.4. Brushing and painting

Another approach for the TF design is to allow the user to select, or probe, the data, not unlike the eyedropper in image editing tools. A TF component is inferred from the point, area, or region of selection. This interaction scheme is also often combined with various forms of automation in the TF generation. This approach can be applied to either the volume rendering view or assisting views, essentially slice views.

8.4.1. Slice view interaction

In medical visualization, a DVR view is often combined with one or more slice views, or multiplanar reconstructions (MPRs). Selecting

a point or region of interest on such an MPR view is then quick and easy, and it also extends to multidimensional data.

Allowing the user to paint on 2D slices of the data can be used as the input for training a machine learning algorithm. The interface proposed by Tzeng *et al.* [TLM03, TLM05] extracts a 10-dimensional attribute space and trains a neural network with the data samples brushed by the user. In their interface, the user can draw strokes for data points to be included or excluded, which is an approach similar to that taken in segmentation interfaces. Soundararajan and Schultz [SS15] use a similar interface in their comparative study of classifiers. The user interface presented by Zhou and Hansen [ZH13] lets users paint line strokes on 2D multivariate slices. Via kernel density estimation, a multivariate TF component is automatically defined and can be further refined by brushing in the volume view or via 2D Gaussians in a reduced 2D space. The system also supports parallel coordinates manipulation. A similar approach has been taken by Zhou and Hansen [ZH14], where advanced selections for more complex areas of interest on slices of multivariate data are performed using a lasso tool that automatically snaps to boundaries. Based on the multivariate attribute space, the selection defines a feature for which a TF definition is automatically generated.

8.4.2. Volume view interaction

Another approach to sample the data is to interact with the volume data directly. Examples are probe based or stroke based as in the work of Zhou and Hansen [ZH14]. However, this approach requires some form of initial simple TF setting to show part of the volume data, which can be a simple gray value gradient.

The probe tool used by Kniss *et al.* [KKH02] takes samples of the volume data on slicing planes displayed in the volume rendering, which in turn result in widgets in the TF domain. The widgets can be manipulated to change the volume rendering interactively. In this way, the user can interact in both the volume domain and the TF domain, referred to as *dual domain interaction*.

An approach for stroke-based TF design has been presented by Ropinski *et al.* [RPSH08], in which the user draws directly on a monochromatic view of the volume and selects features of interest by placing strokes near silhouettes. This approach results in TF components that later can be modified/combined and disabled/enabled to explore the volume. Although Zhou and Hansen [ZH13] primarily painted on 2D slices, the system also allowed the user to further refine the TF definitions by placing strokes or lassos around features of interest, which are then back-projected onto the volume. A very powerful approach for interacting with volume data is the WYSIWYG (What You See is What You Get) approach to TF design presented by Guo *et al.* [GMY11]. They demonstrated a set of intuitive tools to enable interaction directly on the 3D volume rendered image. This work was later extended by Guo and Yuan [GY13] to allow for the design of local transfer functions, achieving expressiveness similar to classical segmentation.

There are, however, considerations to take into account when interacting by drawing on a volume view. Selections in image space correspond to a volumetric region, and thus the actual selection may need to use some thresholding of opacity or a similar operation to identify which voxel values are actually of interest.

8.5. Design galleries

Providing exemplar images from specific TFs is another method for assisting the user in TF generation. Marks *et al.* [MAB*97] introduced the concept of Design Galleries, in which the parameter-space is automatically explored, and thumbnail views are generated to give the user feedback on the results for various parameter settings. This concept applies to many tasks in computer graphics and animation as well as to the exploration of TF settings. The concept has been further explored in various forms by Jankun-Kelly *et al.* [JKM00], who presented a collection of TF settings in the form of a spreadsheet. A two-level interaction approach has been presented by Prauchner *et al.* [PFC05], in which a set of automatically generated TFs are used, and renderings are shown in a gallery style. The approach was also explored by de Moura Pinto and Freitas [PF06]. The user could either fine-tune the selected TF from the gallery or use it as a basis for a stochastic evolutive search in the TF domain. Both papers based the initial TF generation on Kindlemann and Durkin [KD98]. The technique was improved in de Moura Pinto and Freitas [PF08] by providing a history tree to track TF design evolution. The user interface was also improved.

Wu and Qu [WQ07] presented a somewhat different approach but allow the user to select some sample images rendered using predefined TFs. A new TF is then automatically generated by evaluating intermediate results of a mix of source TFs that optimizes the result to match all source images. Guo *et al.* [GLY14] propose the Transfer Function Map, a meta-visualization to organize and utilize expert designed, existing TFs for a given dataset.

A recent approach with the intention to address novice users with no prior knowledge of volume rendering and TFs was presented in Jönsson *et al.* [JFY15]. In this system, a novice user is faced with automatically partitioned ranges and assigned unique colors in a dynamic gallery. The user can then select to adjust the range and opacity and assign a color, using a touch interface and a two-level interaction scheme.

8.6. Higher-level interaction

In this subsection, we discuss work with a specific focus on the workflow and guided interaction. An early taxonomy and rules for selecting colormaps in a general sense were presented in Bergman *et al.* [BRT95] and implemented as a module, named PRAVDAColor, in the IBM Visualization Data Explorer, later OpenDX. This module provides a means to guide the user selecting an appropriate colormap based on the type of data and task at hand.

Castro *et al.* [CKLG98] proposed to add metadata to the TF components, such as component labels. The user can then employ the metadata to specify a TF by selecting components at a higher level, potentially selecting from a list, such as bone, fat, muscle. The authors further describe how components can be saved as presets and automatically adjusted to new datasets. Reapplying existing presets and combinations of TF components could greatly improve the efficiency for the user. In Rezk-Salama *et al.* [RSKK06], semantics is also assigned to the TF components to allow the user to interact with a high-level panel by adjusting the visual appearance of tissues. Selver *et al.* [SFKH07] further explored a complete system of presets and storage of the TF components. An integration of the software

into a DICOM viewer has facilitated the evaluation of the system by medical experts, and overall the TF design process has been shortened.

Since TF specification is often related to domain queries, Natural Language Processing (NLP) may be desired. An interface using Natural Language Processing to allow the user to specify rendering styles in illustrative rendering is proposed in Rautek *et al.* [RBG07] and extended in Rautek *et al.* [RBG08]. The system allows the user to select layered rendering styles with simple expressions, such as "*if principal curvature is not positive then contours are blueish.*"

9. Open Research Questions

While this survey reports on the extensive amount of research conducted on TFs, there remain open questions that call for future study. Many areas of future research are interesting but this section describes some that are of high priority.

Data Modality: TFs are very well suited for specific types of data such as computed tomography. This is not the case for other data modalities. TFs have been less successful for scanning modalities such as MRI, fMRI, or ultrasound. It would be helpful to identify data characteristics that describe which data is suitable for which type of TF. The works described in Section 5.5 can provide starting points in this direction.

Photometric Classification: In almost no case are the physical properties of the TF defined, meaning the material density per unit of physical length of the modeled material. In general, material properties are rarely specified, beyond a basic color and opacity. A few papers deal with styles of the TF. These papers emphasize that the shader or shading model being used can be controlled by the TF. Essentially, two open aspects are given here. As the materials present in the volume have physical properties, the TF should, rather, be defined in relation to these quantities. For example, opacity is derived from material density, which clearly should be defined per unit length. Similarly, the TF should provide the ability to define photometric properties, such as the BRDF or subsurface scattering of the materials to be used. Some research has explored the description of various material properties (see Section 6.2), but such properties are not photometrically based.

Physical Dimensions: Most volume rendering research is established on a unit scale of voxels, and not physical dimensions. Therefore, the result depends on the underlying sampling density. As such, it leaves most relevant measures out of touch with end-users and domain applications, making it difficult to reproduce. For example, finding a feature in a medical dataset, whose size is given like "three voxels wide", is entirely irrelevant. Medical features should have width expressed in a physical unit relating to human physiology. Neglecting the physical background is due to the on purpose simplified TF setup where material classification and visual mapping are intertwined (see Section 3.1). The complex attributes as discussed in Section 4 and Section 5 might serve as an inspiration to more fully consider the physical context of the given application.

Quantifiable Uncertainty: Uncertainty quantification has been an active field of research for quite some time, but, in general, uncertainty representation and visualization has not been solved.

Robust classification algorithms typically produce a probability of the classification. Simply using opacity modulation to represent uncertainty combines these probabilities with the material properties described by the TF. Effective methods for including quantifiable uncertainties or probabilities of classification will benefit the domain users. The work described in Section 4.3 represents initial steps in this promising direction.

Sensitivity, Robustness, Reproducibility: In certain regions of the domain, TFs can be very sensitive, whereas large changes in other regions do not produce large changes in output. Scanned data often contains a certain variability across a slice or within a stack of slices that make up the volume. Classification robustness to such variability has not been well explored. Techniques are needed to cope with these different types of sensitivity. TFs provide the mapping for visual results, but are often not directly quantifiable. Reproducibility would be enhanced with more quantitative results. In this respect, more work on the automated generation of TFs (Section 7) and user interfaces (Section 8) that is aware of sensitivity and robustness is necessary.

Spatial Information: The simplest density-based TFs do not take spatial information into account. Various extensions include local or more global spatial information in the TF specification. For a specific data and task at hand, it would be interesting to identify how much spatial information is needed, from none to all spatial information from the data domain. An open topic is to research and possibly identify the "sweet spot" in going from a TF to a more elaborate classification/segmentation scheme. Various approaches that investigate the wider local neighborhood have been discussed in Sections 4.5, 5.1, 5.4, and 6.4.

User Interfaces: Interacting with TFs remains difficult for end users. As the dimensionality exceeds 2D, such interaction becomes an even greater problem. If we consider the data-user-task design triangle, the user interface could address the data area, or the task area, or be focused on the user. User interfaces for TFs should support the user during both specification and editing. Better interfaces might provide more effective representations of features in either the data range or visual domain, or in combinations thereof. Higher-level interaction techniques, described in Section 8.6, that decrease the cognitive load on the user are a promising area that can bring TFs to the next level of usability.

Scalability: Scalability concerns the size of the data domain, the number of attributes, the number of datasets, etc. Often not a single data volume is investigated but rather entire ensembles or longitudinal time series, perhaps consisting of hundreds of datasets. Examples of work on high-dimensional attribute data is described in several parts of the report, such as in Section 5.3 and Section 7.4. Not only is there variability in a particular dataset, but collections of datasets may have a significant variability between individual datasets. While maintaining quantifiable results, adaptability would benefit the TF specification for such collections. Attributes are increasingly available at each spatial position. Effective techniques for addressing such datasets, perhaps through massive multidimensional TFs, are required. In particular, new techniques may emerge from data analytics that would benefit the TF specification for such data (Section 7).

Transfer Functions for Visualizations in the Wild: The components and setup of a visualization can be easily controlled under laboratory conditions. However, this is not the case in real-world scenarios. Visualization will become increasingly concerned with combining a dynamically changing number of datastreams with limited scope, validity, and accuracy that provide complementing or even contradicting information. The number of streams displayed at the same time may vary considerably. Some of these streams will be volumetric data. How TFs can be adjusted to accommodate this highly volatile situation, and to ensure expressivity while avoiding clutter and visual overload, will be a compelling research question. Starting points for this research could be some of the works in Sections 5, 7.1, and 8.

10. Conclusions

This report has delivered a holistic account of the concept of TFs and systematically described the enabling part TFs play in making volume rendering one of the most successful visualization methods available. The report emphasizes the multifaceted use of the TF, ranging from a classification tool and representation of optical properties, to knowledge encoding and visual design, and as a tool for interactive exploration. We conclude that research on TFs has reached a level of maturity that merits this review of the research contributions made to the field over the past decades. Based on our experience of research in the field, and working closely with domain experts, we have identified areas in which research is needed to take the next important steps to support emerging volumetric exploration concepts and further improve the utility of TF-based data exploration.

Acknowledgments

This research was partially supported by the Department of Energy, National Nuclear Security Administration, under Award Number(s) DE-NA0002375, the DOE SciDAC Institute of Scalable Data Management Analysis and Visualization DOE DE-SC0007446, NIH: Fluorender: An Imaging Tool for Visualization and Analysis of Confocal Data as Applied to Zebrafish Research, R01-GM098151-01, NASA NSSC-NNX16AB93A and NSF ACI-1339881, NSF IIS-1162013. It was further partially supported by the Knut and Alice Wallenberg foundation: Seeing Organ Function KAW.2013.0076. The Swedish Research Council: 2015-05462. Swedish e-Science Research Council (SeRC) and Excellence Center at Linköping - Lund in Information Technology (ELLIIT). This research was made possible in part by the Intel Visual Computing Institute, by the NIH/NCRR Center for Integrative Biomedical Computing, P41-RR12553-10, and by Award Number R01NR14852-01A1 from the National Institute Nursing Research. It was further partially supported by King Abdullah University of Science and Technology (KAUST) baseline funding.

References

- [AD10] ARENS S., DOMIK G.: A survey of transfer functions suitable for volume rendering. In *Eurographics/IEEE VGTC Symposium on Volume Graphics* (2010), VG'10, Eurographics Association, pp. 77–83. doi:10.2312/VG/VG10/077-083. 1
- [AFK*14] AMIRKHANOV A., FRÖHLER B., KASTNER J., GRÖLLER E., HEINZL C.: InSpectr: Multi-modal exploration, visualization, and analysis of spectral data. *Computer Graphics Forum (Proc. of EuroVis)* 33, 3 (June 2014), 91–100. doi:10.1111/cgf.12365. 6
- [AM07] AKIBAY H., MA K.-L.: A tri-space visualization interface for analyzing time-varying multivariate volume data. In *Joint Eurographics - IEEE VGTC Symposium on Visualization* (2007), EuroVis'07, Eurographics Association, pp. 115–122. doi:10.2312/VisSym/EuroVis07/115-122. 16
- [AWHS16] ABBASLOO A., WIENS V., HERMANN M., SCHULTZ T.: Visualizing tensor normal distributions at multiple levels of detail. *IEEE TVCG* 22, 1 (Jan. 2016), 975–984. doi:10.1109/TVCG.2015.2467031. 10
- [BCP*12] BRAMBILLA A., CARNECKY R., PEIKERT R., VIOLA I., HAUSER H.: Illustrative flow visualization: State of the art, trends and challenges. In *Eurographics – STARs* (2012), Cani M.-P., Ganovelli F., (Eds.), Eurographics Association. doi:10.2312/conf/EG2012/stars/075-094. 7
- [BG05] BRUCKNER S., GRÖLLER M. E.: VolumeShop: An interactive system for direct volume illustration. In *IEEE Visualization* (2005), pp. 671–678. doi:10.1109/VISUAL.2005.1532856. 5, 12
- [BG07] BRUCKNER S., GRÖLLER M. E.: Style transfer functions for illustrative volume rendering. *Computer Graphics Forum (Proc. of Eurographics)* 26, 3 (Sept. 2007), 715–724. doi:10.1111/j.1467-8659.2007.01095.x. 5, 11
- [BG09] BRUCKNER S., GRÖLLER M. E.: Instant volume visualization using maximum intensity difference accumulation. *Computer Graphics Forum (Proc. of EuroVis)* 28, 3 (2009), 775–782. doi:10.1111/j.1467-8659.2009.01474.x. 11
- [BGKG05] BRUCKNER S., GRIMM S., KANITSAR A., GRÖLLER M. E.: Illustrative context-preserving volume rendering. In *Joint Eurographics - IEEE VGTC Symposium on Visualization* (2005), EuroVis'05, Eurographics Association, pp. 69–76. doi:10.2312/VisSym/EuroVis05/069-076. 12
- [BGKG06] BRUCKNER S., GRIMM S., KANITSAR A., GRÖLLER M. E.: Illustrative context-preserving exploration of volume data. *IEEE TVCG* 12, 6 (2006), 1559–1569. doi:10.1109/TVCG.2006.96. 12
- [BRB*13] BRAMON R., RUIZ M., BARDERA A., BOADA I., FEIXAS M., SBERT M.: Information theory-based automatic multimodal transfer function design. *IEEE Journal of Biomedical and Health Informatics* 17, 4 (July 2013), 870–880. doi:10.1109/JBHI.2013.2263227. 14
- [BRT95] BERGMAN L., ROGOWITZ B., TREINISH L.: A rule-based tool for assisting colormap selection. In *IEEE Visualization* (1995), pp. 118–125, 444. doi:10.1109/VISUAL.1995.480803. 17
- [BVMG08] BALABANIAN J.-P., VIOLA I., MÖLLER T., GRÖLLER M. E.: Temporal styles for time-varying volume data, 2008. Poster presented at 3D Data Processing, Visualization, and Transmission. URL: <http://www.cg.tuwien.ac.at/research/publications/2008/balabanian-2008-tst/>. 12
- [BZGV14] BISTA S., ZHUO J., GULLAPALLI R. P., VARSHNEY A.: Visualization of brain microstructure through spherical harmonics illumination of high fidelity spatio-angular fields. *IEEE TVCG (Proc. of Vis.)* 20, 12 (Dec. 2014), 2516–2525. doi:10.1109/TVCG.2014.2346411. 10
- [CBW*12] CHILDS H., BRUGGER E., WHITLOCK B., MEREDITH J., AHERN S., PUGMIRE D., BIAGAS K., MILLER M., HARRISON C., WEBER G. H., KRISHNAN H., FOGAL T., SANDERSON A., GARTH C., BETHEL E. W., CAMP D., RÜBEL O., DURANT M., FAVRE J. M., NÁVRATIL P.: VisIt: An end-user tool for visualizing and analyzing very large data. In *High Performance Visualization—Enabling Extreme-Scale Scientific Insight*, Chapman & Hall/CRC Computational Science. CRC Press, Oct 2012, pp. 357–372. 5
- [CIB15] CIBC: 2015. ImageVis3D: An interactive visualization software system for large-scale volume data. Scientific Computing and Imaging Institute (SCI). URL: <http://www.imagevis3d.org>. 5

- [CKLG98] CASTRO S., KÖNIG A., LÖFFELMANN H., GRÖLLER E.: Transfer function specification for the visualization of medical data, 1998. Technical Report: TR-186-2-98-12. URL: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.48.7705.5.13.14.15.17>
- [CM08] CORREA C., MA K.-L.: Size-based transfer functions: A new volume exploration technique. *IEEE TVCG (Proc. of Vis.)* 14, 6 (Nov. 2008), 1380–1387. doi:10.1109/TVCG.2008.162. 7
- [CM09a] CORREA C., MA K. L.: The occlusion spectrum for volume classification and visualization. *IEEE TVCG (Proc. of Vis.)* 15, 6 (Nov. 2009), 1465–1472. doi:10.1109/TVCG.2009.189. 7, 8
- [CM09b] CORREA C., MA K.-L.: Visibility-driven transfer functions. In *IEEE Pacific Visualization Symposium* (2009), PacificVis 2009, pp. 177–184. doi:10.1109/PACIFICVIS.2009.4906854. 14
- [CM11] CORREA C. D., MA K.-L.: Visibility histograms and visibility-driven transfer functions. *IEEE TVCG* 17, 2 (Feb. 2011), 192–204. doi:10.1109/TVCG.2010.35. 14, 15
- [CMH*01] CSÉBFAI B., MROZ L., HAUSER H., KÖNIG A., GRÖLLER E.: Fast visualization of object contours by non-photorealistic volume rendering. *Computer Graphics Forum (Proc. of Eurographics)* 20, 3 (2001), 452–460. doi:10.1111/1467-8659.00538. 11, 12
- [CMPS97] CIGNONI P., MONTANI C., PUPPO E., SCOPIGNO R.: Multiresolution representation and visualization of volume data. *IEEE TVCG* 3, 4 (Oct. 1997), 352–369. doi:10.1109/2945.646238. 15
- [CNCO15] CAI L.-L., NGUYEN B. P., CHUI C.-K., ONG S.-H.: Rule-Enhanced Transfer Function Generation for Medical Volume Visualization. *Computer Graphics Forum (Proc. of EuroVis)*, 3 (2015). doi:10.1111/cgcf.12624. 9
- [CR08] CABAN J., RHEINGANS P.: Texture-based transfer functions for direct volume rendering. *IEEE TVCG (Proc. of Vis.)* 14, 6 (Nov. 2008), 1364–1371. doi:10.1109/TVCG.2008.169. 8
- [DANS10] DANIELS JR J., ANDERSON E. W., NONATO L. G., SILVA C. T.: Interactive vector field feature identification. *IEEE TVCG (Proc. of Vis.)* 16, 6 (2010), 1560–8. doi:10.1109/TVCG.2010.170. 7
- [DCH88] DREBIN R. A., CARPENTER L., HANRAHAN P.: Volume rendering. In *SIGGRAPH* (1988), ACM, pp. 65–74. doi:10.1145/54852.378484. 1, 3, 4
- [DEP05] DESGRANGES P., ENGEL K., PALADINI G.: Gradient-free shading: A new method for realistic interactive volume rendering. In *International Workshop on Vision, Modeling and Visualization* (2005), VMV'05, pp. 209–216. 5
- [dMPF07] DE MOURA PINTO F., FREITAS C. M. D. S.: Design of multi-dimensional transfer functions using dimensional reduction. In *Joint Eurographics - IEEE VGTC Symposium on Visualization* (2007), EuroVis'07, Eurographics Association, pp. 131–138. doi:10.2312/VisSym/EuroVis07/131-138. 10, 15
- [EHK*06] ENGEL K., HADWIGER M., KNISS J., REZK-SALAMA C., WEISKOPF D.: *Real-time volume graphics*. CRC Press, 2006. 2, 3
- [EMRY02] EBERT D. S., MORRIS C. J., RHEINGANS P., YOO T. S.: Designing effective transfer functions for volume rendering from photographic volumes. *IEEE TVCG* 8, 2 (Apr. 2002), 183–197. doi:10.1109/2945.998670. 6
- [FAT99] FUJISHIRO I., AZUMA T., TAKESHIMA Y.: Automating transfer function design for comprehensible volume rendering based on 3D field topology analysis. In *IEEE Visualization* (1999), pp. 467–563. doi:10.1109/VISUAL.1999.809932. 10, 14
- [FBT98] FANG S., BIDDLECOME T., TUCERYAN M.: Image-based transfer function design for data exploration in volume visualization. In *IEEE Visualization* (1998), pp. 319–326. doi:10.1109/VISUAL.1998.745319. 14
- [FHG*09] FRITZ L., HADWIGER M., GEIER G., PITTINO G., GRÖLLER M. E.: A visual approach to efficient analysis and quantification of ductile iron and reinforced sprayed concrete. *IEEE TVCG (Proc. of Vis.)* 15, 6 (2009), 1343–1350. doi:10.1109/TVCG.2009.115. 7, 8, 12
- [FTAT00] FUJISHIRO I., TAKESHIMA Y., AZUMA T., TAKAHASHI S.: Volume data mining using 3D field topology analysis. *Computer Graphics and Applications* 20, 5 (Sept. 2000), 46–51. doi:10.1109/38.865879. 10, 14
- [FW08] FALK M., WEISKOPF D.: Output-sensitive 3D line integral convolution. *IEEE TVCG* 14, 4 (July 2008), 820–834. doi:10.1109/TVCG.2008.25. 7
- [GLY14] GUO H., LI W., YUAN X.: Transfer function map. In *IEEE Pacific Visualization Symposium* (2014), PacificVis 2014, pp. 262–266. doi:10.1109/PacificVis.2014.24. 17
- [GMY11] GUO H., MAO N., YUAN X.: WYSIWYG (what you see is what you get) volume visualization. *IEEE TVCG (Proc. of Vis.)* 17, 12 (Dec. 2011), 2106–2114. doi:10.1109/TVCG.2011.261. 17
- [GXY11] GUO H., XIAO H., YUAN X.: Multi-dimensional transfer function design based on flexible dimension projection embedded in parallel coordinates. In *IEEE Pacific Visualization Symposium* (2011), PacificVis 2011, pp. 19–26. doi:10.1109/PACIFICVIS.2011.5742368. 16
- [GXY12] GUO H., XIAO H., YUAN X.: Scalable multivariate volume visualization and analysis based on dimension projection and parallel coordinates. *IEEE TVCG* 18, 9 (Sept. 2012), 1397–1410. doi:10.1109/TVCG.2012.80. 16
- [GY13] GUO H., YUAN X.: Local WYSIWYG volume visualization. In *IEEE Pacific Visualization Symposium* (2013), PacificVis 2013, pp. 65–72. doi:10.1109/PacificVis.2013.6596129. 17
- [HA04] HENDERSON A., AHRENS J.: *The Paraview guide: a parallel visualization application*. Kitware, Inc., New York, 2004. URL: <http://opac.inria.fr/record=b1117983.5>
- [HBG11] HAIDACHER M., BRUCKNER S., GRÖLLER E.: Volume analysis using multimodal surface similarity. *IEEE TVCG (Proc. of Vis.)* 17, 12 (2011), 1969–1978. doi:10.1109/TVCG.2011.258. 7
- [HBH03] HADWIGER M., BERGER C., HAUSER H.: High-quality two-level volume rendering of segmented data sets on consumer graphics hardware. In *IEEE Visualization* (2003), pp. 301–308. doi:10.1109/VISUAL.2003.1250386. 5, 12
- [HBKG08] HAIDACHER M., BRUCKNER S., KANITSAR A., GRÖLLER M. E.: Information-based transfer functions for multimodal visualization. In *Eurographics Workshop on Visual Computing for Biomedicine* (2008), VCBM 2008, Eurographics Association, pp. 101–108. doi:10.2312/VCBM/VCBM08/101-108. 10
- [HHKP96] HE T., HONG L., KAUFMAN A., PFISTER H.: Generation of transfer functions with stochastic search techniques. In *IEEE Visualization* (1996), pp. 227–ff. doi:10.1109/Visual.1996.568113. 14
- [HKG00] HLADUVKA J., KÖNIG A., GRÖLLER E.: Curvature-based transfer functions for direct volume rendering. In *Spring Conference on Computer Graphics* (2000), SCCG 2000, pp. 58–65. URL: <https://www.cg.tuwien.ac.at/research/vis/vismed/CurvatureTF/CurvatureTF.pdf>. 7
- [HLY07] HERNELL F., LJUNG P., YNNERMAN A.: Efficient ambient and emissive tissue illumination using local occlusion in multiresolution volume rendering. In *Eurographics/IEEE VGTC Symposium on Volume Graphics* (2007), Hege H.-C., Machiraju R., Möller T., Sramek M., (Eds.), VG'07, Eurographics Association. doi:10.2312/VG/VG07/001-008. 11
- [HPB*10] HAIDACHER M., PATEL D., BRUCKNER S., KANITSAR A., GRÖLLER E.: Volume visualization based on statistical transfer-function spaces. In *IEEE Pacific Visualization Symposium* (2010), PacificVis 2010, pp. 17–24. doi:10.1109/PACIFICVIS.2010.5429615. 6, 7
- [HSS*05] HADWIGER M., SIGG C., SCHARSACH H., BÜHLER K., GROSS M.: Real-time ray-casting and advanced shading of discrete isosurfaces. *Computer Graphics Forum (Proc. of Eurographics)* 24, 3 (2005), 303–312. 7
- [IVJ12] IP C. Y., VARSHNEY A., JÁJÁ J.: Hierarchical exploration

- of volumes using multilevel segmentation of the intensity-gradient histograms. *IEEE TVCG (Proc. of Vis.)* 18, 12 (2012), 2355–2363. doi:10.1109/TVCG.2012.231. 8, 14
- [JFY15] JÖNSSON D., FALK M., YNNERMAN A.: Intuitive Exploration of Volumetric Data Using Dynamic Galleries. *IEEE TVCG (Proc. of Vis.)* 22, 1 (2015), 896–905. doi:10.1109/TVCG.2015.2467294. 17
- [JH09] JOHNSON C., HUANG J.: Distribution-driven visualization of volume data. *IEEE TVCG* 15, 5 (Sept. 2009), 734–746. doi:10.1109/TVCG.2009.25. 9
- [JKM00] JANKUN-KELLY T., MA K.-L.: A spreadsheet interface for visualization exploration. In *IEEE Visualization* (2000), pp. 69–76. doi:10.1109/VISUAL.2000.885678. 17
- [KB08] KRAUS M., BÜRGER K.: Interpolating and downsampling RGBA volume data. In *International Workshop on Vision, Modeling and Visualization* (2008), pp. 323–332. 13
- [KBKG08] KOHLMANN P., BRUCKNER S., KANITSAR A., GRÖLLER E.: LiveSync++: enhancements of an interaction metaphor. In *Graphics Interface Conference* (2008), ACM, pp. 81–88. doi:10.1145/1375714.1375729. 13
- [KD98] KINDLMANN G., DURKIN J. W.: Semi-automatic generation of transfer functions for direct volume rendering. In *IEEE Symposium on Volume Visualization* (1998), pp. 79–86. doi:10.1109/SVV.1998.729588. 13, 14, 17
- [KG01] KÖNIG A., GRÖLLER E.: Mastering transfer function specification by using VolumePro technology. In *Spring Conference on Computer Graphics* (2001), vol. 17, pp. 279–286. URL: <https://www.cg.tuwien.ac.at/research/publications/2000/Koenig-2000-ATFS/TR-186-2-00-07Paper.pdf>, doi:10.1.1.43.5954. 15
- [KKH02] KNISS J., KINDLMANN G., HANSEN C.: Multidimensional transfer functions for interactive volume rendering. *IEEE TVCG* 8, 3 (July 2002), 270–285. doi:10.1109/TVCG.2002.1021579. 5, 6, 15, 16, 17
- [KPB12] KROES T., POST F. H., BOTHA C. P.: Exposure render: An interactive photo-realistic volume rendering framework. *PloS one* 7, 7 (2012). doi:10.1371/journal.pone.0038586. 13
- [KSC*10] KIM H. S., SCHULZE J. P., CONE A. C., SOSINSKY G. E., MARTONE M. E.: Dimensionality reduction on multi-dimensional transfer functions for multi-channel volume data sets. *Information Visualization* 9, 3 (June 2010), 167–180. doi:10.1057/ivs.2010.6. 10
- [KSW06] KRÜGER J., SCHNEIDER J., WESTERMANN R.: ClearView: An interactive context preserving hotspot visualization technique. *IEEE TVCG (Proc. of Vis.)* 12, 5 (Sept. 2006), 941–948. doi:10.1109/TVCG.2006.124. 7, 12
- [KUS*05] KNISS J. M., UITERT R. V., STEPHENS A., LI G. S., TADIZEN T., HANSEN C.: Statistically quantitative volume visualization. In *IEEE Visualization* (2005), pp. 287–294. doi:10.1109/VISUAL.2005.1532807. 3, 7, 9
- [KVH84] KAJIYA J. T., VON HERZEN B. P.: Ray tracing volume densities. *Computer Graphics (Proc. of SIGGRAPH)* (Jan. 1984), 165–174. doi:10.1145/800031.808594. 1, 3
- [KW99] KINDLMANN G., WEINSTEIN D.: Hue-balls and lit-tensors for direct volume rendering of diffusion tensor fields. In *IEEE Visualization* (1999), pp. 183–524. doi:10.1109/VISUAL.1999.809886. 10
- [KWH00] KINDLMANN G., WEINSTEIN D., HART D.: Strategies for direct volume rendering of diffusion tensor fields. *IEEE TVCG* 6, 2 (Apr. 2000), 124–138. doi:10.1109/2945.856994. 10
- [KWTM03] KINDLMANN G., WHITAKER R., TADIZEN T., MÖLLER T.: Curvature-based transfer functions for direct volume rendering: Methods and applications. In *IEEE Visualization* (2003), pp. 513–520. doi:10.1109/VISUAL.2003.1250414. 7
- [LE05] LU A., EBERT D. S.: Example-based volume illustrations. In *IEEE Visualization* (2005), pp. 655–662. doi:10.1109/VISUAL.2005.1532854. 11
- [Lev88] LEVOY M.: Display of surfaces from volume data. *Computer Graphics and Applications* 8, 3 (May 1988), 29–37. doi:10.1109/38.511. 1, 3, 5
- [LJHY14] LINDHOLM S., JÖNSSON D., HANSEN C., YNNERMAN A.: Boundary aware reconstruction of scalar fields. *IEEE TVCG (Proc. of Vis.)* 20, 12 (2014), 2447–2455. doi:10.1109/TVCG.2014.2346351. 12, 13
- [LLL*10] LINDHOLM S., LJUNG P., LUNDSTROM C., PERSSON A., YNNERMAN A.: Spatial conditioning of transfer functions using local material distributions. *IEEE TVCG (Proc. of Vis.)* 16, 6 (Nov. 2010), 1301–1310. doi:10.1109/TVCG.2010.195. 4, 9
- [LLPY07] LUNDSTROM C., LJUNG P., PERSSON A., YNNERMAN A.: Uncertainty visualization in medical volume rendering using probabilistic animation. *IEEE TVCG (Proc. of Vis.)* 13, 6 (Nov. 2007), 1648–1655. doi:10.1109/TVCG.2007.70518. 3, 7, 12
- [LLY05] LUNDSTRÖM C., LJUNG P., YNNERMAN A.: Extending and simplifying transfer function design in medical volume rendering using local histograms. In *Joint Eurographics - IEEE VGTC Symposium on Visualization* (2005), EuroVis'05, Eurographics Association, pp. 263–270. doi:10.2312/VisSym/EuroVis05/263-270. 8
- [LLY06a] LUNDSTRÖM C., LJUNG P., YNNERMAN A.: Local histograms for design of transfer functions in direct volume rendering. *IEEE TVCG* 12, 6 (2006), 1570–1579. doi:10.1109/TVCG.2006.100. 5, 8, 9
- [LLY06b] LUNDSTRÖM C., LJUNG P., YNNERMAN A.: Multi-dimensional transfer function design using sorted histograms. In *Eurographics/IEEE VGTC Workshop on Volume Graphics* (2006), Machiraju R., Möller T., (Eds.), VG'06, Eurographics Association. doi:10.2312/VG/VG06/001-008. 16
- [LS16] LIU X., SHEN H. W.: Association analysis for visual exploration of multivariate scientific data sets. *IEEE TVCG* 22, 1 (Jan. 2016), 955–964. doi:10.1109/TVCG.2015.2467431. 16
- [LWT*14] LIU S., WANG B., THIAGARAJAN J. J., BREMER P. T., PASCUCCI V.: Multivariate volume visualization through dynamic projections. In *IEEE Symposium on Large Data Analysis and Visualization* (2014), LDAV 2014, pp. 35–42. doi:10.1109/LDAV.2014.7013202. 14
- [LYL*06] LUNDSTRÖM C., YNNERMAN A., LJUNG P., PERSSON A., KNUTSSON H.: The α -histogram: Using spatial coherence to enhance histograms and transfer function design. In *Joint Eurographics - IEEE VGTC Symposium on Visualization* (2006), EuroVis'06, Eurographics Association, pp. 227–234. doi:10.2312/VisSym/EuroVis06/227-234. 9, 15
- [LZY*07a] LI J., ZHOU L., YU H., LIANG H., WANG L.: Classification for volume rendering of industrial CT based on moment of histogram. In *IEEE Conference on Industrial Electronics and Applications* (2007), pp. 913–918. ICIEA 2007. doi:10.1109/ICIEA.2007.4318542. 8
- [LZY*07b] LI J., ZHOU L., YU H., LIANG H., ZHANG L.: Classification for volume rendering of industrial CT based on minimum cross entropy. In *IEEE International Conference on Mechatronics and Automation* (2007), pp. 2710–2715. ICMA 2007. doi:10.1109/ICMA.2007.4303986. 5, 8
- [MAB*97] MARKS J., ANDALMAN B., BEARDSLEY P. A., FREEMAN W., GIBSON S., HODGINS J., KANG T., MIRTICH B., PFISTER H., RUMMLER W., RYALL K., SEIMS J., SHIEBER S.: Design galleries: A general approach to setting parameters for computer graphics and animation. In *SIGGRAPH* (1997), ACM, pp. 389–400. doi:10.1145/258734.258887. 14, 17
- [Max95] MAX N.: Optical models for direct volume rendering. *IEEE TVCG* 1, 2 (June 1995), 99–108. doi:10.1109/2945.468400. 2
- [ME02] MORRIS C. J., EBERT D.: Direct volume rendering of photographic volumes using multi-dimensional color-based transfer functions. In *Joint Eurographics - IEEE TVCG Symposium on Visualization* (2002), VisSym 2002, Eurographics Association, pp. 115–ff. doi:10.2312/VisSym/VisSym02/115-124. 6

- [MJW*13] MACIEJEWSKI R., JANG Y., WOO I., JÄNICKE H., GAITHER K. P., EBERT D. S.: Abstracting attribute space for transfer function exploration and design. *IEEE TVCG* 19, 1 (Jan. 2013), 94–107. doi:10.1109/TVCG.2012.105. 15
- [MNKT01] MURAKI S., NAKAI T., KITA Y., TSUDA K.: An attempt for coloring multichannel MR imaging data. *IEEE TVCG* 7, 3 (July 2001), 265–274. doi:10.1109/2945.942694. 6, 10
- [MW09] MANKE F., WÜNSCHE B.: Texture-enhanced direct volume rendering. In *International Conference on Computer Graphics Theory and Applications* (2009), VISIGRAPP 2009, SciTePress, pp. 185–190. doi:10.5220/0001772701850190. 11
- [MWCE09] MACIEJEWSKI R., WOO I., CHEN W., EBERT D. S.: Structuring feature space: A non-parametric method for volumetric transfer function generation. *IEEE TVCG* 15, 6 (2009), 1473–1480. doi:10.1109/TVCG.2009.185. 8, 14
- [PBL*04] PARK S., BUDGE B., LINSSEN L., HAMANN B., JOY K.: Multi-dimensional transfer functions for interactive 3D flow visualization. In *Pacific Conference on Computer Graphics and Applications* (2004), PCCGA'04, IEEE, pp. 177–185. doi:10.1109/PCCGA.2004.1348348. 7
- [PF06] PINTO F., FREITAS C.: Two-level interaction transfer function design combining boundary emphasis, manual specification and evolutive generation. In *Brazilian Symposium on Computer Graphics and Image Processing* (2006), SIBGRAPI 2006, IEEE, pp. 281–288. doi:10.1109/SIBGRAPI.2006.45. 17
- [PF08] PINTO F. D. M., FREITAS C. M. D. S.: Volume visualization and exploration through flexible transfer function design. *Computers & Graphics* 32, 5 (Oct. 2008), 540–549. doi:10.1016/j.cag.2008.08.006. 17
- [PFC05] PRAUCHNER J., FREITAS C., COMBA J.: Two-level interaction approach for transfer function specification. In *Brazilian Symposium on Computer Graphics and Image Processing* (2005), SIBGRAPI 2005, IEEE, pp. 265–272. doi:10.1109/SIBGRAPI.2005.52. 14, 17
- [PGT*08] PATEL D., GIERTSEN C., THURMOND J., GJELBERG J., GRÖLLER M. E.: The seismic analyzer: Interpreting and illustrating 2d seismic data. *IEEE TVCG* 14, 6 (2008), 1571–1578. doi:10.1109/TVCG.2008.170. 12
- [PGTG07] PATEL D., GIERTSEN C., THURMOND J., GRÖLLER E.: Illustrative rendering of seismic data. In *International Workshop on Vision, Modeling and Visualization* (2007), VMV'07, pp. 13–22. URL: http://www.cg.tuwien.ac.at/research/publications/2007/patel_daniel_2007_IRSD/. 12
- [PHBG09] PATEL D., HAIDACHER M., BALABANIAN J., GRÖLLER M. E.: Moment curves. In *IEEE Pacific Visualization Symposium* (2009), PacificVis 2009, pp. 201–208. doi:10.1109/PACIFICVIS.2009.4906857. 9
- [PLB*01] PFISTER H., LORENSEN B., BAJAJ C., KINDLMANN G., SCHROEDER W., MACHIRAJU R., LEE J.: The transfer function bake-off. *Computer Graphics and Applications* 21, 3 (May 2001), 16–22. doi:10.1109/38.920623. 1, 4
- [PM04] POTTS S., MÖLLER T.: Transfer functions on a logarithmic scale for volume rendering. In *Graphics Interface Conference* (2004), GI'04, Canadian Human-Computer Communications Society, pp. 57–63. URL: <http://dl.acm.org/citation.cfm?id=1006058.1006066>. doi:10.1145/1006058.1006066. 15
- [PRMH10] PRASSNI J.-S., ROPINSKI T., MENSMAAN J., HINRICHS K. H.: Shape-based transfer functions for volume visualization. In *IEEE Pacific Visualization Symposium* (2010), PacificVis 2010, pp. 9–16. doi:10.1109/PacificVis.2010.5429624. 8, 10, 14
- [PRW11] PFAFFELMOSER T., REITINGER M., WESTERMANN R.: Visualizing the positional and geometrical variability of isosurfaces in uncertain scalar fields. *Computer Graphics Forum (Proc. of EuroVis)* 30, 3 (2011), 951–960. doi:10.1111/j.1467-8659.2011.01944.x. 7, 14
- [RBG07] RAUTEK P., BRUCKNER S., GRÖLLER E.: Semantic layers for illustrative volume rendering. *IEEE TVCG (Proc. of Vis.)* 13, 6 (Nov. 2007), 1336–1343. doi:10.1109/TVCG.2007.70591. 11, 18
- [RBG08] RAUTEK P., BRUCKNER S., GRÖLLER M. E.: Interaction-dependent semantics for illustrative volume rendering. *Computer Graphics Forum (Proc. of EuroVis)* 27, 3 (May 2008), 847–854. doi:10.1111/j.1467-8659.2008.01216.x. 11, 18
- [RBS05] ROETTGER S., BAUER M., STAMMINGER M.: Spatialized transfer functions. In *Joint Eurographics - IEEE VGTC Symposium on Visualization* (2005), EuroVis'05, Eurographics Association, pp. 271–278. doi:10.2312/VisSym/EuroVis05/271-278. 14
- [RPSH08] ROPINSKI T., PRASSNI J., STEINICKE F., HINRICHS K.: Stroke-based transfer function design. In *Eurographics/IEEE VGTC Symposium on Volume and Point-Based Graphics* (2008), VGPBG'08, Eurographics Association, pp. 41–48. doi:10.2312/VG/VG-PBG08/041-048. 17
- [RSHSG00] REZK-SALAMA C., HASTREITER P., SCHERER J., GREINER G.: Automatic adjustment of transfer functions for 3D volume visualization. In *International Workshop on Vision, Modeling and Visualization* (2000), VMV 2000, pp. 357–364. URL: http://www.cg.informatik.uni-siegen.de/data/Publications/2000/rezk_Publ.2000.11.pdf. 13
- [RSKK06] REZK-SALAMA C., KELLER M., KOHLMANN P.: High-level user interfaces for transfer function design with semantics. *IEEE TVCG (Proc. of Vis.)* 12, 5 (Sept. 2006), 1021–1028. doi:10.1109/TVCG.2006.148. 13, 16, 17
- [ŠBSG06] ŠEREDA P., BARTOLI A. V., SERLIE I. W. O., GERRITSEN F. A.: Visualization of boundaries in volumetric data sets using lh histograms. *IEEE TVCG* 12, 2 (Mar. 2006), 208–218. doi:10.1109/TVCG.2006.39. 9, 14
- [SCO98] SHIBOLET O., COHEN-OR D.: Coloring voxel-based objects for virtual endoscopy. In *IEEE Symposium on Volume Visualization* (1998), pp. 15–22. doi:10.1109/SVV.1998.729580. 11
- [Sel15] SELVER M. A.: Exploring brushlet based 3D textures in transfer function specification for direct volume rendering of abdominal organs. *IEEE TVCG* 21, 2 (Feb. 2015), 174–187. doi:10.1109/TVCG.2014.2359462. 10
- [SFKH07] SELVER M. A., FISCHER F., KUNTALP M., HILLEN W.: A software tool for interactive generation, representation, and systematic storage of transfer functions for 3D medical images. *Computer Methods and Programs in Biomedicine* 86, 3 (June 2007), 270–280. doi:10.1016/j.cmpb.2007.03.008. 17
- [SG09] SELVER M. A., GUZELIS C.: Semiautomatic transfer function initialization for abdominal visualization using self-generating hierarchical radial basis function networks. *IEEE TVCG* 15, 3 (May 2009), 395–409. doi:10.1109/TVCG.2008.198. 6, 14
- [SJEG05] SVAKHINE N. A., JANG Y., EBERT D., GAITHER K.: Illustration and photography inspired visualization of flows and volumes. In *IEEE Visualization* (2005), pp. 687–694. doi:10.1109/VISUAL.2005.1532858. 7
- [SKMH14] SICAT R., KRÜGER J., MÖLLER T., HADWIGER M.: Sparse PDF volumes for consistent multi-resolution volume rendering. *IEEE TVCG (Proc. of Vis.)* 20, 12 (Dec. 2014), 2417–2426. doi:10.1109/TVCG.2014.2346324. 13
- [SLMSC13] SRINIVASAN S., LJUNG P., MCDERMOTT B., SMITH-CASEM M.: Adaptive volume rendering for ultrasound color flow diagnostic imaging, Apr. 2013. US Patent 8,425,422. URL: <http://www.google.com/patents/US8425422>. 5
- [SS15] SOUNDARARAJAN K. P., SCHULTZ T.: Learning probabilistic transfer functions: A comparative study of classifiers. *Computer Graphics Forum (Proc. of EuroVis)* (2015). doi:10.1111/cgf.12623. 7, 15, 17
- [STF*03] SERLIE I., TRUYEN R., FLORIE J., POST F., VLIET L., VOS F.: *MICCAI*. Springer, 2003, ch. Computed Cleansing for Virtual

- Colonoscopy Using a Three-Material Transition Model, pp. 175–183. doi:10.1007/978-3-540-39903-2_22. 9
- [ŠVG06] ŠEREDA P., VILANOVA A., GERRITSEN F. A.: Automating transfer function design for volume rendering using hierarchical clustering of material boundaries. In *Joint Eurographics - IEEE VGTC Symposium on Visualization* (2006), Santos B. S., Ertl T., Joy K., (Eds.), EuroVis'06, Eurographics Association. doi:10.2312/VisSym/EuroVis06/243-250. 9, 14
- [TLB*11] THOMPSON D., LEVINE J. A., BENNETT J. C., BREMER P.-T., GYULASSY A., PASCUCCHI V., PÉBAY P. P.: Analysis of large-scale scalar data using hixels. In *IEEE Symposium on Large Data Analysis and Visualization* (2011), LDAV 2011, pp. 23–30. doi:10.1109/LDAV.2011.6092313. 13
- [TLM03] TZENG F.-Y., LUM E. B., MA K.-L.: A novel interface for higher-dimensional classification of volume data. In *IEEE Visualization* (2003), pp. 66–. doi:10.1109/VISUAL.2003.1250413. 17
- [TLM05] TZENG F.-Y., LUM E., MA K.-L.: An intelligent system approach to higher-dimensional classification of volume data. *IEEE TVCG* 11, 3 (May 2005), 273–284. doi:10.1109/TVCG.2005.38. 14, 15, 17
- [TLM02] TAKANASHI I., LUM E., MA K.-L., MURAKI S.: ISpace: interactive volume data classification techniques using independent component analysis. In *Pacific Conference on Computer Graphics and Applications* (2002), PCCGA 2002, IEEE, pp. 366–374. doi:10.1109/PCCGA.2002.1167880. 6
- [TM04] TZENG F.-Y., MA K.-L.: A cluster-space visual interface for arbitrary dimensional classification of volume data. In *Joint Eurographics - IEEE TVCG Symposium on Visualization* (2004), VisSym 2004, Eurographics Association, pp. 17–24. doi:10.2312/VisSym/VisSym04/017-024. 8, 9
- [TPD06] TAPPENBECK A., PREIM B., DICKEN V.: Distance-based transfer function design: Specification methods and applications. In *Simulation and Visualization* (2006), pp. 259–274. 7
- [TTF04] TAKAHASHI S., TAKESHIMA Y., FUJISHIRO I.: Topological volume skeletonization and its application to transfer function design. *Graphical Models* 66, 1 (2004), 24 – 49. doi:10.1016/j.gmod.2003.08.002. 10
- [TTFN05] TAKESHIMA Y., TAKAHASHI S., FUJISHIRO I., NIELSON G. M.: Introducing topological attributes for objective-based visualization of simulated datasets. In *Eurographics/IEEE VGTC Workshop on Volume Graphics* (2005), VG'05, Eurographics Association. doi:10.2312/VG/VG05/137-145. 10
- [VKG04] VIOLA I., KANITSAR A., GRÖLLER E.: Importance-driven volume rendering. In *IEEE Visualization* (2004), pp. 139–146. doi:10.1109/VISUAL.2004.48. 11
- [VKG05] VIOLA I., KANITSAR A., GRÖLLER M. E.: Importance-driven feature enhancement in volume visualization. *IEEE TVCG* 11, 4 (2005), 408–418. doi:10.1109/TVCG.2005.62. 11
- [WCLC06] WANG L., CHEN X., LI S., CAI X.: General adaptive transfer functions design for volume rendering by using neural networks. In *International Conference on Neural Information Processing - Volume Part II* (2006), ICONIP'06, Springer-Verlag, pp. 661–670. doi:10.1007/11893257_74. 15
- [WCZ*11] WANG Y., CHEN W., ZHANG J., DONG T., SHAN G., CHI X.: Efficient volume exploration using the Gaussian mixture model. *IEEE TVCG* 17, 11 (2011), 1560–1573. doi:10.1109/TVCG.2011.97. 8
- [WDC*07] WEBER G. H., DILLARD S. E., CARR H., PASCUCCHI V., HAMANN B.: Topology-controlled volume rendering. *IEEE TVCG* 13, 2 (Mar. 2007), 330–341. doi:10.1109/TVCG.2007.47. 10
- [WK12] WANG L., KAUFMAN A.: Importance driven automatic color design for direct volume rendering. *Computer Graphics Forum (Proc. of EuroVis)* 31, 3 (2012), 1305 – 1314. doi:10.1111/j.1467-8659.2012.03123.x. 14
- [WQ07] WU Y., QU H.: Interactive transfer function design based on editing direct volume rendered images. *IEEE TVCG* 13, 5 (Sept. 2007), 1027–1040. doi:10.1109/TVCG.2007.1051. 17
- [WWT09] WONG H.-C., WONG U.-H., TANG Z.: Direct volume rendering by transfer function morphing. In *International Conference on Information, Communications and Signal Processing* (2009), ICICS'09, IEEE, pp. 857–860. doi:10.1109/ICICS.2009.5397587. 12
- [WZK12] WANG L., ZHAO X., KAUFMAN A.: Modified dendrogram of attribute space for multidimensional transfer function design. *IEEE TVCG* 18, 1 (Jan. 2012), 121–131. doi:10.1109/TVCG.2011.23. 8
- [WZL*12] WANG Y., ZHANG J., LEHMANN D. J., THEISEL H., CHI X.: Automating transfer function design with valley cell-based clustering of 2D density plots. *Computer Graphics Forum (Proc. of EuroVis)* 31, 3 (June 2012), 1295–1304. doi:10.1111/j.1467-8659.2012.03122.x. 14
- [XTY*11] XIANG D., TIAN J., YANG F., YANG Q., ZHANG X., LI Q., LIU X.: Skeleton cuts — an efficient segmentation method for volume rendering. *IEEE TVCG* 17, 9 (Sept. 2011), 1295–1306. doi:10.1109/TVCG.2010.239. 10
- [YMC06] YOUNESY H., MÖLLER T., CARR H.: Improving the quality of multi-resolution volume rendering. In *Joint Eurographics - IEEE VGTC Symposium on Visualization* (2006), EuroVis'06, Eurographics Association. doi:10.2312/VisSym/EuroVis06/251-258. 12
- [YNCP05] YUAN X., NGUYEN M. Z., CHEN B., PORTER D. H.: High dynamic range volume visualization. In *IEEE Visualization* (2005), pp. 327–334. doi:10.1109/VISUAL.2005.1532812. 13, 15
- [YNCP06] YUAN X., NGUYEN M. X., CHEN B., PORTER D. H.: HDR VolVis: high dynamic range volume visualization. *IEEE TVCG* 12, 4 (July 2006), 433–445. doi:10.1109/TVCG.2006.72. 13, 15
- [ZH13] ZHOU L., HANSEN C. D.: Transfer function design based on user selected samples for intuitive multivariate volume exploration. In *IEEE Pacific Visualization Symposium* (2013), PacificVis 2013, IEEE, pp. 73–80. doi:10.1109/PacificVis.2013.6596130. 10, 17
- [ZH14] ZHOU L., HANSEN C.: GuideME: Slice-guided semiautomatic multivariate exploration of volumes. *Computer Graphics Forum (Proc. of EuroVis)*, 3 (2014). doi:10.1111/cgf.12371. 10, 17
- [ZK10] ZHAO X., KAUFMAN A.: Multi-dimensional reduction and transfer function design using parallel coordinates. In *Eurographics/IEEE VGTC Symposium on Volume Graphics* (2010), VG'10, Eurographics Association, pp. 69–76. doi:10.2312/VG/VG10/069-076. 10, 16
- [ZSH12] ZHOU L., SCHOTT M., HANSEN C.: Technical section: Transfer function combinations. *Computers & Graphics* 36, 6 (Oct. 2012), 596–606. doi:10.1016/j.cag.2012.02.007. 6, 16
- [ZT09] ZHOU J., TAKATSUKA M.: Automatic transfer function generation using contour tree controlled residue flow model and color harmonics. *IEEE TVCG (Proc. of Vis.)* 15, 6 (Nov. 2009), 1481–1488. doi:10.1109/TVCG.2009.120. 10