

Predicate-based Focus-and-Context Visualization for 3D Ultrasound

Christian Schulte zu Berge, Maximilian Baust, Ankur Kapoor, and Nassir Navab

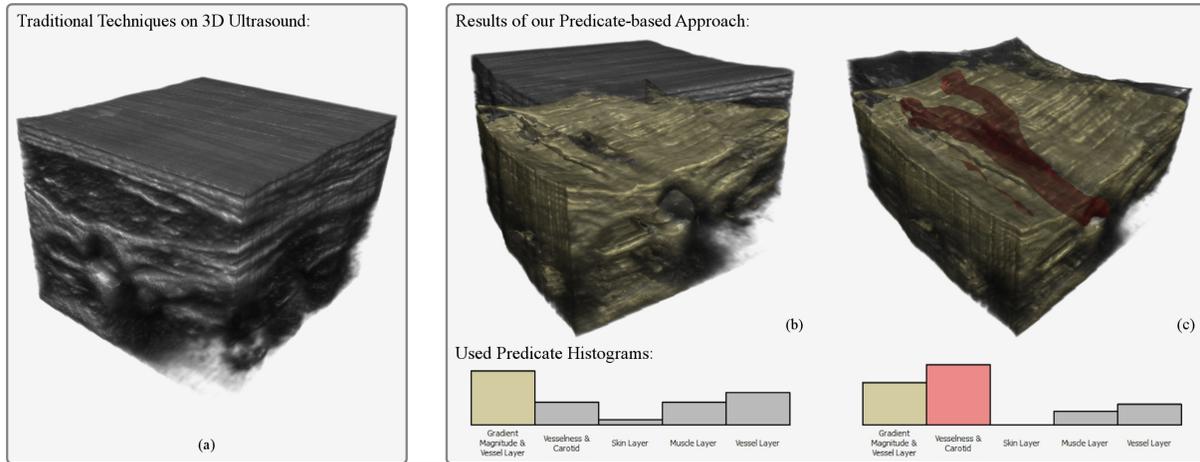


Fig. 1: Our predicate-based classification approach allows for more insightful visualization of 3D ultrasound volumes. Here, traditional transfer function-based techniques suffer from occlusion of inner target anatomy such as the carotid artery (a). Facilitated through the predicate histogram, our technique is capable of yielding a focus-and-context visualization of not only the surrounding tissue (b), but also the target anatomy itself (c).

Abstract—Direct volume visualization techniques offer powerful insight into volumetric medical images and are part of the clinical routine for many applications. Up to now, however, their use is mostly limited to tomographic imaging modalities such as CT or MRI. With very few exceptions, such as fetal ultrasound, classic volume rendering using one-dimensional intensity-based transfer functions fails to yield satisfying results in case of ultrasound volumes. This is particularly due its gradient-like nature, a high amount of noise and speckle, and the fact that individual tissue types are rather characterized by a similar texture than by similar intensity values. Therefore, clinicians still prefer to look at 2D slices extracted from the ultrasound volume. In this work, we present an entirely novel approach to the classification and compositing stage of the volume rendering pipeline, specifically designed for use with ultrasonic images. We introduce point predicates as a generic formulation for integrating the evaluation of not only low-level information like local intensity or gradient, but also of high-level information, such as non-local image features or even anatomical models. Thus, we can successfully filter clinically relevant from non-relevant information. In order to effectively reduce the potentially high dimensionality of the predicate configuration space, we propose the predicate histogram as an intuitive user interface. This is augmented by a scribble technique to provide a comfortable metaphor for selecting predicates of interest. Assigning importance factors to the predicates allows for focus-and-context visualization that ensures to always show important (focus) regions of the data while maintaining as much context information as possible. Our method naturally integrates into standard ray casting algorithms and yields superior results in comparison to traditional methods in terms of visualizing a specific target anatomy in ultrasound volumes.

Index Terms—Direct Volume Rendering, Ultrasound, Classification, Predicate Function, User Interface

1 INTRODUCTION

Direct volume visualization techniques such as volume ray casting are today’s state-of-the-art algorithms for the visualization of three-dimensional medical images. In order to approximate the physics of light transport, most techniques use global transfer functions for the classification step of the classic volume rendering pipeline. While they

yield impressive results for images from tomographic reconstructions such as computed tomography (CT) or magnetic resonance imaging (MRI), global 1D transfer functions reach their limits when being applied to 3D ultrasound images. Though they have recently become part of clinical practice due to the advances in today’s 3D freehand ultrasound systems, effective and generic visualization tools are still missing for this imaging modality. This is mainly because several characteristics of B-mode ultrasound (brightness mode ultrasound where pixel intensities relate to the change of acoustic impedance and thus the reflectance of the sound wave) let its volumetric visualization suffer from severe occlusion artifacts as shown in Figure 2.

- Christian Schulte zu Berge, Maximilian Baust, and Nassir Navab are with the Chair for Computer Aided Medical Procedures, Technische Universität München, Germany, E-mail: christian.szb@in.tum.de.
- Ankur Kapoor is with Imaging and Computer Vision, Siemens Corporation, Corporate Technology, Princeton, NJ, USA, E-mail: ankur.kapoor@siemens.com.

Manuscript received 31 Mar. 2014; accepted 1 Aug. 2014; date of publication xx xxx 2014; date of current version xx xxx 2014.

For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.

In lack of better alternatives, clinicians prefer to look at 2D slices extracted from 3D ultrasound volumes. Such multi-planar reconstructions (MPRs) allow for the visualization of arbitrary planes, which is not directly possible with 2D ultrasound. However, MPRs are hardly able to expose spatial context and connectivity information to the clinician. Rendering extracted geometry from the image obtained by seg-

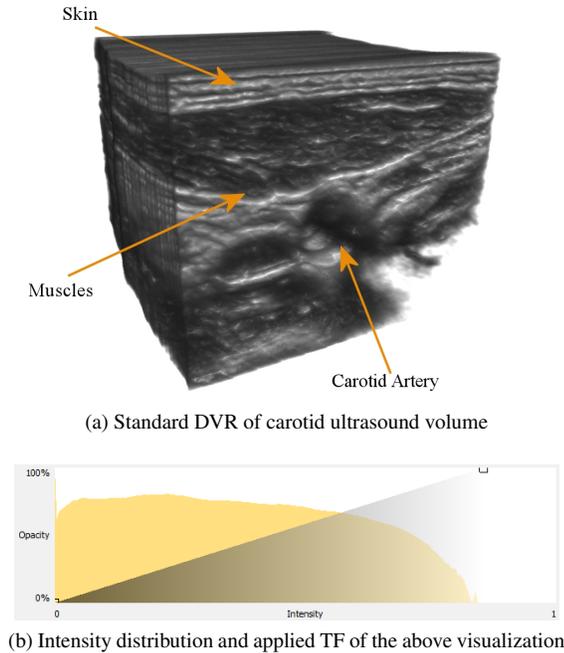


Fig. 2: **Occlusion artifacts** of traditional methods: (a) While one can see the outer shell of the different layers, clinical relevant information such as the path and shape of the carotid artery and its bifurcation are hidden. (b) Due to the homogeneous distribution of ultrasound intensities and their highly context-sensitive nature, there is no chance in discriminating relevant features through their intensities.

mentation algorithms, as an alternative, may show connectivity information but does no longer expose the original image data (i.e. ultrasound intensities) and in particular the speckle texture, which is an essential part of ultrasound imaging. Therefore, this kind of indirect visualization is not well suited for many clinical routines.

In order to achieve clinically helpful direct volume visualization, we propose a completely novel approach to the classification stage together with an intuitive tool for setting up the rendering result: We introduce the concept of point predicates, which evaluate both local and global features of the ultrasound image and are defined on every sampling point. This facilitates to perform classification also based on high-level non-local information such as speckle or texture or even anatomical models/segmentations. By further annotating each predicate with an importance factor, we can naturally implement relevance-based visualization ensuring that important anatomies are always visible in the rendering, while preserving context information where possible in order to show better spatial clues.

Despite the recent advances regarding quality, interactivity, and usability, volumetric visualization still has not been fully accepted by most clinicians for their workflow and is, if at all, only used for publications or patient presentations. Even with CT imaging, where the image intensities (Hounsfield units) directly correspond to a physical property that allows for direct discrimination of tissue types, most radiologists still prefer to scroll through the stack of 2D slices and make up the 3D model in their minds instead of looking at 3D visualizations. This is mainly due to the lack of usability of today's classification approaches, where changing the volume rendering to show different anatomy relates to adaption of the transfer function. The limited expressiveness and intuitiveness, as well as the high dimensionality of the parameter domain, in particular with multi-dimensional transfer functions, makes their setup a tedious and cumbersome task that many clinicians have difficulties with [20].

With this issue in mind, we further introduce the predicate histogram as an effective tool for reducing the dimensionality of the predicate configuration space and facilitating its manipulation. Together

with the descriptive semantics of predicates, it provides the user with easy and intuitive interaction with the point predicates to setup the rendering. This user interface was designed to allow for interactive exploration of clinically relevant information and switching between visualizations of different target anatomies with minimal efforts in a highly intuitive way. This is further enriched by a scribble technique providing a painting metaphor to specify classification directly in the image domain.

2 RELATED WORK

Our proposed method has three important properties:

1. It is particularly designed for the visualization of ultrasound volumes.
2. It integrates different levels of information, ranging from low-level local image intensities to even anatomical models, into a single consistent formulation that is exposed to the user using an intuitive widget.
3. It exploits illustrative focus-and-context rendering techniques.

Since the body of literature on volume rendering techniques is large, we focus this section on the most closely related works regarding the above topics and refer the interested reader to the book of Preim and Botha [16] for a more exhaustive overview.

2.1 Visualization of Ultrasound Volumes

Compared to tomographic imaging modalities such as CT or MRI, where direct volume rendering can show very distinct visualizations of the anatomy, B-mode ultrasound images provide special challenges to classification causing these techniques to fail in yielding helpful visualizations. Since ultrasound imaging exploits the echo generated by density changes at tissue interfaces of different acoustic impedance, it rather shows the changes in physical properties than the physical properties themselves. As a consequence, ultrasound images are fundamentally different to those obtained from tomographic imaging modalities. In addition to this gradient-like nature, ultrasound is a highly directional modality and suffers from a considerable amount of noise. These properties prohibit the straightforward application of standard classification techniques for direct volume visualization (cf. Figure 2).

One of the few methods, besides the early work of Sakas et al. [19], specifically targeting these challenges for visualization, is the work of Fattal and Lischinski [5], who propose a variational approach to opacity classification that allows to extract smooth surfaces from 3D ultrasound volumes. However, their work is mostly attributed to fetal ultrasound, since it basically shows only a single surface and does not allow for blending of multiple layers. Mann et al. propose a volumetric ultrasound system augmenting B-mode ultrasound intensities with elasticity information using two-dimensional transfer functions in order to yield more distinct visualizations [14].

2.2 Multi-Dimensional Classification Schemes

While exploiting additional information during classification and therefore introducing multi-dimensional transfer functions [11, 23, 15] (for instance based on post-processing) is a viable way towards useful 3D ultrasound visualization, it significantly increases the parameter domain to setup the rendering. Designing effective user interfaces for transfer function and rendering setup is an extensive and important field of research, as especially non-expert users have difficulties with mapping the complex parameter domain to semantic features for visualization. Therefore, Rezk-Salama et al. propose using principal component analysis to map a small set of semantic parameters to the potentially large transfer function parameter space and claim that this can be learned from clinicians [20]. Wang et al. introduce the Modified Dendrogram as a means for mapping high dimensional transfer functions into 2D space to facilitate their setup for the user [27]. The resulting user interface, however, still has a high amount of non-descriptive and thus unintuitive parameters. A different approach

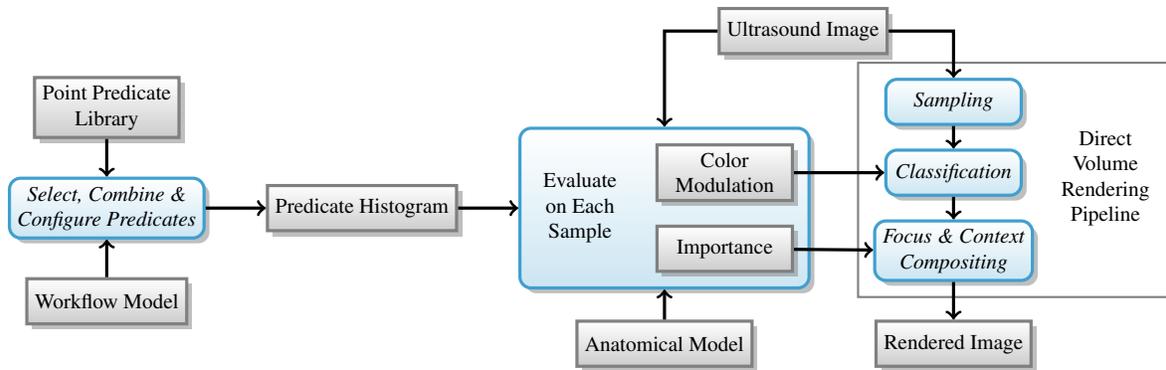


Fig. 3: **Schematic diagram** of the proposed predicate-based rendering technique: The user selects a set of descriptive point predicates from the library and configures their settings, importance factors and color modulation using the intuitive predicate histogram widget. These predicates are then evaluated on each sample during classification and eventually accumulated using relevance-based compositing exploiting the importance factors. While the predicates may evaluate both local and global features in the ultrasound image as well as external information (e.g. an anatomical model), the application-specific selection of the predicates can potentially be learned and supported with a workflow model.

towards solving the classification problem by introducing a semantic layer is the work of Rautek et al. [17], who implement a fuzzy logic evaluation of the semantic descriptions on the GPU and combine it with interaction-dependent rendering. While they achieve impressive relevance-based rendering results, their work is mostly focused on integrating user interaction into the visualization and does not address the special challenges of ultrasonic images.

2.3 Focus-and-Context Visualization

Illustrative relevance-based visualization is an efficient technique to tackle occlusion problems in volume visualization. They extend the classic direct volume rendering compositing scheme towards focus-and-context rendering that defines a focus region, which must not be occluded, and context regions that are shown with less priority for better spatial understanding. Approaches not relying on transfer functions, such as the Importance-Based Accumulated Transparency Modulation introduced by Wan and Hansen [26] extend the Maximum Intensity Difference Accumulation (MIDA) compositing scheme to allow relevance-based visualization without the need to setup a transfer function for classification. The work of Bruckner et al. allows for context-preserving see-through rendering by evaluating local shading information with two global parameters, which however do not have a direct semantic meaning [2]. The ClearView technique of Krüger et al. exploits curvature information as well as distance metrics to determine sample importances and introduces different shading and compositing techniques to map relevance to optical properties [13]. De Moura Pinto and Freitas introduce a further importance-aware compositing scheme, which is mathematically motivated and justified and integrates very well with the standard direct volume rendering pipeline [3].

2.4 Our Approach

Our approach is inspired by line predicates used in flow visualization, where streamline tracking yields an extensive number of streamlines that represent the global connectivity of the data very well, but at the same time it greatly suffer from occlusion of important features. Here, line predicates offer an effective technique to filter the flow field for certain features such as vortices or high-velocity jets [21, 1, 12]. However, since line predicates are applied to geometry representations and simply toggle streamline visibility, they are very limited to this specific application. Our approach can thus be seen as a generalization of this idea providing a consistent formulation for relevance-based rendering with particular focus and application on 3D ultrasound visualization.

3 METHODS

Our predicate-based approach is designed to fully integrate into a standard volume rendering pipeline and consists of three steps, as illus-

trated in the schematic diagram of Figure 3. We define a point predicate P as a boolean-valued function f_P on the image domain X augmented with an importance factor κ_P and a color modulation δ_P :

$$P := (f_P : X \rightarrow \{\text{true}, \text{false}\}, \kappa_P, \delta_P). \quad (1)$$

After selecting the predicates to apply from the point predicate library, the user specifies the rendering outcome by configuring κ_P and δ_P . This process is heavily supported by the predicate histogram, which we propose as a user interface for rendering setup. During the classification stage of the rendering, the ray caster applies each predicate to each sampling point yielding a color modulation for each sample, which are eventually accumulated using a focus-and-context compositing technique exploiting the predicates' importance factors.

The term point predicate suggests that the predicate can be evaluated at every point within the image domain X . However, it is essential that predicates are not limited to local values but also to features of local environments or even of global nature, such as texture information or anatomical models and segmentations. Furthermore, we do not pose any presumptions on the spatial representation of X .

3.1 Predicate Library

Our current implementation consists of a point predicate library based on a variety of methods to evaluate both local and global features in the image.

Range-based predicates apply to intervals of scalar measures in the image such as intensity and gradient magnitude. While their sole expressiveness is rather limited, prone to inter-data set variability, and in particular does not go beyond traditional 1D/2D transfer functions, they are an essential part in combination with the other predicates. For instance, since high curvature regions often carry important visual depth and context cues, assigning high importance values to the gradient magnitude predicates may yield significant visual improvements.

Due to the direction dependency of ultrasound images, *direction-based predicates* are an important factor for our predicate-based 3D ultrasound visualization. They exploit additional information on the ultrasound scanning direction, which we annotate to the compounded ultrasound volume. The gradient angle predicate evaluates the scanline direction in a local context by computing the angle between scanline and the smoothed gradient. Since the reflection of the ultrasound wave depends on the incident angle with the interface, this ultrasound specific point predicate is a powerful tool to highlight or mask certain structures.

As a third group we implemented a variety of predicates based on derived measures. Since they apply a configurable threshold to the derived measure, they can be seen as extension to the range-based predicates.

The *Signal-to-Noise ratio predicate* computes the variance-based SNR of the ultrasound image in a local neighborhood of the sample position and can be used to select regions with certain entropy. We therefore build for each voxel the weighted variance within a uniform $7 \times 7 \times 7$ kernel and perform windowing on the results. The windowing parameters are user-controlled in our current implementation, since an elaborate evaluation is beyond the scope of this work.

Furthermore, we implemented a GPU version of Frangi’s vesselness filter [6] and bound its results to a *vesselness predicate*. By computing local Hessians for each voxel on multi-level Gaussians of the ultrasound volume and performing an eigenvalue analysis, the Frangi filter provides a scalar measure of presence of tubular structures. Therefore, this predicate is particularly useful in vascular imaging.

The presence of ultrasound confidence maps [9] can be exploited with the *confidence predicate*. Confidence Maps compute a relative per-pixel attenuation of the ultrasound signal using a random-walks approach. This attenuation could be interpreted as uncertainty information and therefore directly be integrated into the volume rendering pipeline. Our experiments however show only limited use of the additional information, which may be an issue with the confidence maps themselves. Nevertheless, we think that this is a promising direction for future work to integrate uncertainty visualization into our pipeline since future work on Confidence Maps, such as [8], may improve their quality and validity.

Finally, we implemented a *label predicate* that evaluates on optionally existing labeling information of the image in two different ways. For our evaluation in Section 6 we performed a manual segmentation on the data sets partitioning the ultrasound volumes into semantic layers of anatomy such as skin, muscle and bone. This was performed in the style of Surface Function Actives segmentation [4], which could yield a similar result in a (semi-)automatic fashion [10]. The segmentation layers were stored as B-spline surfaces in voxel space. With these, the label predicate can then directly determine to which layer the current sample belongs. The second version of the label predicate uses a 3D voxel grid as input data, where each voxel stores a bit mask with the labeling information. With the cost of discretization, this predicate allows the same evaluation of segmentation information but in a more generic way. This yields a very natural way to seamlessly integrate segmentations into the ray casting process and offers a powerful tool to accurately define importance factors or color modulations for specific regions.

As our work focuses on providing a novel approach together with a high-level tool, the presented point predicate library is certainly not complete, but rather provides a proof-of-concept implementation. Future work by the community may yield a variety of further application-dependent predicates that evaluate both local and global features in the image to yield even better classification results.

3.2 Predicate Combination

To allow sufficient generality, we keep the individual predicates simple with each just evaluating a single feature in the image domain. We achieve further flexibility by assembling multiple predicates P_1, \dots, P_n into *meta predicates* using combination operators. Our boolean formulation of the predicate function therefore yields the classic boolean operators:

NOT: $\overline{P_1}$ yields the negation of the original predicate and allows to semantically complement the accepted domain of the original predicate.

AND: $P_1 \wedge \dots \wedge P_n$ yields a predicate that satisfies the constraints of all predicates P_1 through P_n . Often it is desired to have this operation in order to semantically restrict the accepted domain to the intersection of the accepted domains of each predicate, as also shown in Figure 4 where the second and the third predicate from the left are of this type.

OR: $P_1 \vee \dots \vee P_n$ yields a predicate that satisfies the constraints of either of the predicates P_1 through P_n . Since our formulation computes the final sample importance by computing a weighted

sum of the predicates’ importances (cf. Section 3.4), the OR operator can also be obtained by suitable setup of the importance factors.

3.3 Predicate Selection & Setup

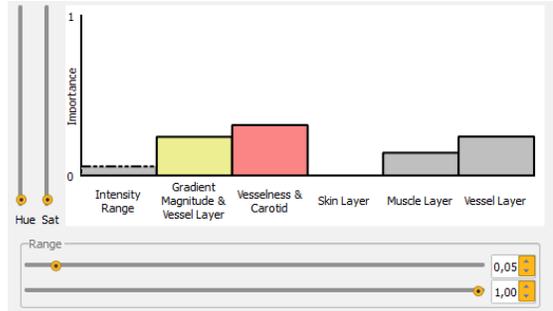


Fig. 4: Our **proposed widget** shows the predicate histogram and allows an easy and intuitive configuration of the point predicates for relevance-based rendering: The main area shows the distribution of predicate importances and enables easy manipulation using drag and drop. Two sets of controls allow for further manipulation of the currently selected predicate: The sliders on the left set up the color modulation; the controls on the bottom set up optional predicate parameters.

The point predicates to apply are selected from the user or, alternatively, given by a workflow model later potentially learned specifically for the given anatomy and application. Each predicate has a small set of parameters, defining the predicate importance κ and an optional color modulation δ having components hue $\delta^{(H)}$ and saturation $\delta^{(S)}$. Certain predicates may have further optional parameters defining where the predicate yields true (e.g. intensity range).

We propose to combine predicate selection and configuration via the *predicate histogram*, a single intuitive widget as shown in Figure 4. The heights of the bars directly represent the relative importance weighting of the predicates, while their fill color shows the color modulation. The user can manipulate the importance of each predicate with intuitive drag and drop interaction on the histogram bar and will directly see the results of his/her actions, as the predicate-based rendering is evaluated in real-time. In order to further reduce the parameter space without losing flexibility, the predicate histogram normalizes the sum of the predicates weights:

$$\sum_j \kappa_j = 1. \quad (2)$$

Hence, increasing the importance of one predicate automatically decreases the importances of the other predicates and vice versa.

We place two additional sets of manipulators around the predicate histogram widget, which allow controlling additional parameters of the currently selected predicate: Two vertical sliders on the left allow the user to configure the color modulation in terms of hue and saturation. The resulting color is visualized in the predicate histogram as fill color of the corresponding bar. Furthermore, widgets to control feature specific predicate parameters (such as the range where the predicate yields true) are placed below the histogram in a horizontal layout.

The predicate selection process is currently a fully manual task of explorative nature. However, we argue that achieving satisfying rendering results is significantly more intuitive and easier than with traditional 1D/2D transfer functions, since the parameter space of our point predicate technique is much more expressive due to its semantic nature and the whole process runs in real-time providing interactive feedback. Furthermore, the predicate histogram reduces the parameter space effectively in size through the normalization of predicate importances (cf. Equation (2)).

3.4 Classification & Compositing

For classification and compositing, our predicate-based ray caster evaluates the selected predicates on each input sample point s to compute the sample importance $\kappa(s)$ and sample color modulation $\delta(s)$. As we wish the individual importances to contribute in an additive manner, we define $\kappa(s)$ as

$$\kappa(s) = \frac{\sum_{j=1}^n \chi_j(s) I(\kappa_j)}{\sum_{j=1}^n \chi_j(s)}, \quad (3)$$

where $I(\kappa_j)$ denotes the impact function for predicate j and $\chi_j(s)$ the characteristic function for predicate j (i.e. indicating whether $f_j(s)$ is true). We would like the following two conditions to hold:

1. A non-linear amplification of importance differences, meaning that $\kappa_i \sim \kappa_j \Rightarrow I(\kappa_i) \sim I(\kappa_j)$ but $\kappa_i < \kappa_j \Rightarrow I(\kappa_i) \ll I(\kappa_j)$. This allows for better usability since the user does not need to be pixel-perfect when assigning the same importance to multiple predicates. This condition can be rewritten to

$$\kappa_i - \kappa_j < I(\kappa_i) - I(\kappa_j), \quad (4)$$

which holds for all differentiable I with $I'(x) > 1$.

2. In the case of all predicates having the same importance factor assigned, we wish $\kappa(s)$ to be always 1:

$$\forall i, j : \kappa_i = \kappa_j \Rightarrow \forall s : \kappa(s) = 1. \quad (5)$$

Experimentally, we found $I(\kappa_j) := (n \cdot \kappa_j)^2$, where n is the total number of predicates, to yield good results and satisfy both conditions. This gives us the following full definition (cf. Equation (3)):

$$\kappa(s) = \frac{\sum_{j=1}^n \chi_j(s) \cdot (n \cdot \kappa_j)^2}{\sum_{j=1}^n \chi_j(s)}. \quad (6)$$

It should be noted that one certainly can select very small κ_i, κ_j , so that $n \cdot \kappa_i < 0.5$ and $n \cdot \kappa_j < 0.5$ violate the first condition. However, in such cases there is at least one κ_k with $n \cdot \kappa_k > 1$ due to Equation (2). We argue that this effect is even desirable as predicate k shall have the major impact on the visualization and we can neglect the importance difference between predicates i and j .

In order to specify the optical properties of the sample s , we compute the sample color modulation in terms of hue $\delta^{(H)}(s)$ and saturation $\delta^{(S)}(s)$ in HSL color space. The sample luminance as well as its opacity are set to the original ultrasound intensity. This serves two goals: It reduces the dimensionality of the parameter space, and even more importantly allows to retain the appearance of the original ultrasound data, which we see as an essential part of ultrasound visualization. The color modulation is computed in an additive manner using the predicate contributions

$$\omega_j(s) := \chi_j(s) \cdot (n \cdot \kappa_j)^2 \quad (7)$$

as weights. The hue modulation is additionally weighted by the saturation modulation to avoid undesirable effects in cases where large hue shifts with small saturation modulations are combined with small hue shifts with large saturation modulations.

$$\delta^{(S)}(s) = \frac{1}{\sum_{j=1}^n \omega_j(s)} \cdot \sum_{j=1}^n \omega_j(s) \cdot \delta_j^{(S)},$$

$$\delta^{(H)}(s) = \frac{1}{\sum_{j=1}^n \omega_j(s) \delta_j^{(S)}} \cdot \sum_{j=1}^n \omega_j(s) \delta_j^{(S)} \cdot \delta_j^{(H)}. \quad (8)$$

With sample importance, color and opacity set, the samples are fed into a focus and context compositing scheme. While our proposed method should work with most relevance-based compositing techniques, we use the one proposed by De Moura Pinto and Freitas [3], which has various desirable properties:

1. Its well motivated front-to-back recurrence scheme formulation allows straightforward integration into standard ray casters.
2. In the absence of importance values or with all samples being equally important, it resolves to the standard emission-absorption-based volume rendering integral.
3. Its input parameters are clear and limited to a scalar importance value besides the standard optical properties of color and opacity. We therefore map $\kappa(s)$ to I_s , $\delta(s)$ to C_s and $I(s)$ to α_s as introduced in [3].

The full recurrence scheme for compositing can be found in the Appendix.

4 EXTENSIONS

Our technique integrates very well into the existing volume rendering pipeline. This is also valid for a variety of possible extensions that current state-of-the-art methods provide to facilitate standard transfer function-based classification, such as data driven techniques to generate transfer function presets, shape-based transfer functions, or image-driven transfer function setup where the user can manipulate the classification by direct interaction with the original image. Many of these works can be easily transferred to our predicate-based approach for direct volume rendering. As an example, we present how to integrate a scribble technique into our framework similar to the interface proposed by Tzeng et al. [25] and the stroke-based transfer function design by Ropinski et al. [18].

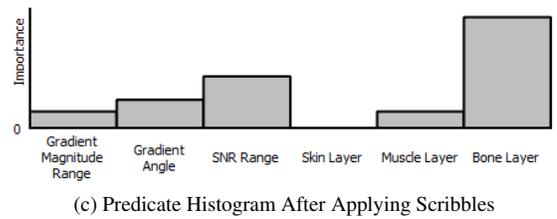
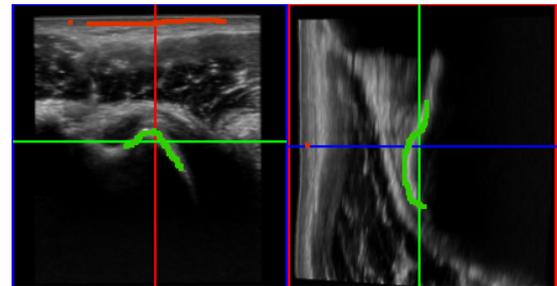
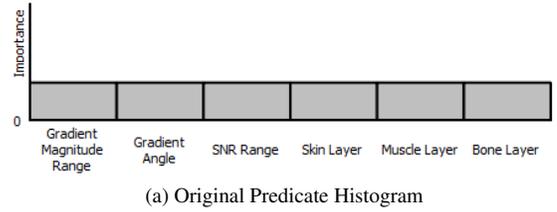


Fig. 5: Illustration of scribble-based Predicate Histogram Setup: (a) shows the original predicate histogram. (b) By delineating image features to show (green scribble along bone structure) and image features to hide (red scribble in skin layer) in the cross-sectional MPRs, the user can configure the predicate histogram directly in the image domain. (c) shows the updated predicate histogram after applying the scribbles.

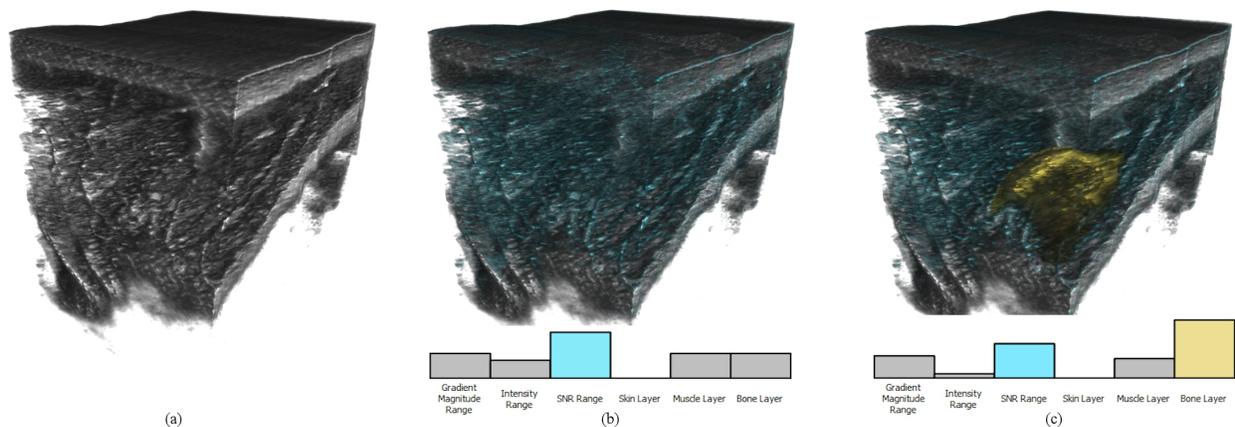


Fig. 6: Visualization results for the **shoulder data set** together with the used predicate histograms: Compared to the standard transfer function-based visualization (a), (b) was generated by exploiting the SNR Range predicate to mark fibrous structures in the muscle layer. The final rendering (c) additionally highlights the bone surface as focus region.

4.1 Scribble-based Predicate Histogram Setup

While the predicate histogram widget already provides an intuitive metaphor to setup the predicate configuration, it still requires the user to perform an implicit mapping between the semantic meaning of each predicate and its corresponding features in the image. To further facilitate the predicate importance setup, we implemented a scribble technique that offers the user the possibility to directly interact with the ultrasound image by using a painting metaphor to mark focus and context regions in cross-sectional slice views of the volume as illustrated in Figure 5.

The scribbles yield a set of positive samples that should be emphasized in the volume rendering and a set of negative samples that should be less prominent. Since emphasizing certain regions relates to increasing the importances of the corresponding predicates and since decreasing importances results in decreased visibility, we transform the scribbles into a modulation of the predicate importances κ_i .

For each predicate i we denote the number of positively labeled samples where the predicate yields true by N_i^+ and the number of negatively labeled samples where the predicate yields true by N_i^- . We compute an importance modulation vector m as

$$m_i = \frac{q(N_i^+ - N_i^-)}{\sum_j N_j^+ - N_j^-}, \quad (9)$$

where q describes the percentage by which the current predicate histogram is altered through the scribbles, which we empirically set to 0.25. In order to keep the predicate histogram in its normalized state (Equation (2)), we first normalize m so that $\sum_j m_j = 0$ before we update each importance κ_i by adding m_i .

This approach can certainly be extended with more elaborate importance modulation or even automatic predicate parameter configuration. However, since this is beyond the scope of this paper we consider it as future work.

5 IMPLEMENTATION

We implemented our predicate-based ray casting technique entirely using OpenGL 4 and GLSL, so that apart from setting up the predicate logic, all evaluations and computations are performed on the GPU using a single fragment shader. Furthermore, our formulation is designed to seamlessly integrate into the standard direct volume rendering pipeline.

5.1 Predicate Selection & Setup

The client-side selection and setup is automatically transformed into shader code, for which we use a building block like system similar to the shader templates in [17]. Every point predicate provides three blocks:

1. A *GLSL header* defining uniform names for the predicate parameters and further optional definitions.
2. A *closed GLSL expression* defining how to evaluate the predicate with respect to the input data.
3. A *function to setup the shader*, which is called by the renderer (prior to rendering) and assigns the parameter values to the corresponding uniforms.

The classification scheme from Equations (6) and (8) is transformed into a single GLSL function combining the evaluation expressions and yielding the final sample intensity and color modulation. Together with the predicates' GLSL headers, the predicate evaluation function is injected into the renderers fragment shader. Meta predicates are easily implemented using a composite pattern.

Potentially expensive recompilation of the shader has to be performed only when the set of selected point predicates changes, as this is the single event changing the definition of uniform variables and predicate expressions. Subsequent rendering setup by configuring the predicates' importance factors, color modulations, and optional parameters is entirely implemented by updating the uniform values, which is even less overhead than an update of the transfer function texture as performed by traditional transfer function-based renderers.

To implement the proposed scribble technique, we need to encode for each pixel which predicates yield true. This is easily done by generating a 3D volume where each voxel encodes the results of the predicates' characteristic functions in a bit mask.

5.2 Integration into the Standard DVR Pipeline

Integrating our predicate-based classification and compositing into a standard GLSL ray caster is straightforward and needs only minimal changes in the shader code: Instead of a texture lookup on the transfer function, the shader gathers all necessary input data for the predicates and calls the predicate evaluation function yielding final sample importance and color modulation. The color modulation is then applied to sample intensity in HSL space, yielding the sample color value and opacity. Together with the sample importance they are directly fed into the compositing scheme as defined in [3].

6 RESULTS

For evaluation we acquired ultrasound sweeps of three distinct anatomies (shoulder, carotid artery, and achilles tendon) using an ACUSON S2000TM ultrasound machine equipped with an Acuson 9L4 linear transducer and electromagnetic tracking hardware. These sweeps were reconstructed using an advanced backward compounding algorithm exploiting frame orientation information to yield high quality volumes. Further details on this algorithm can be found in [22].

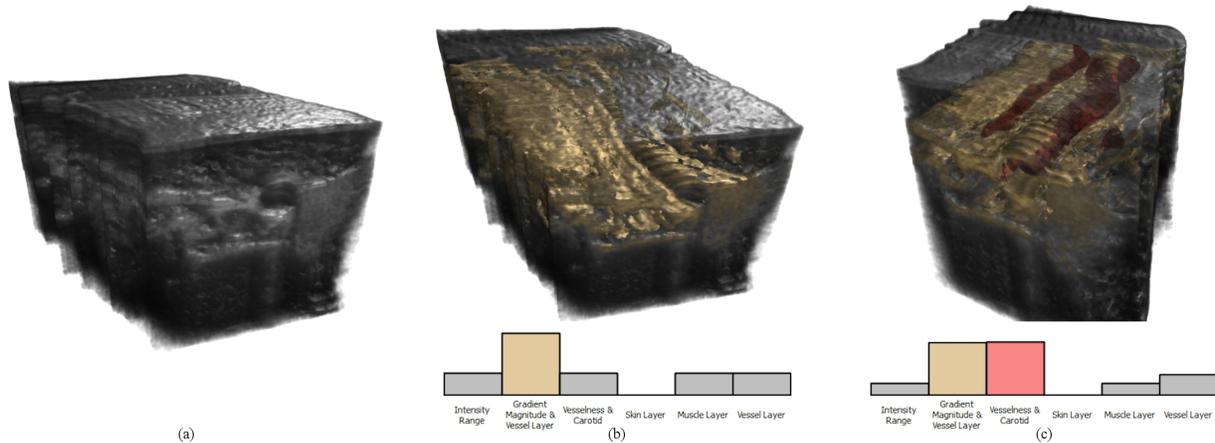


Fig. 7: Visualization results for the second **carotid artery data set** together with the used predicate histograms: Compared to the standard transfer function-based visualization (a), the predicate setup in (b) removes the skin layer and highlights vascular structures. The vesselness predicate in (c) allows to further show the vessel path (the bifurcation is only partly present in this data set and thus can not be seen).

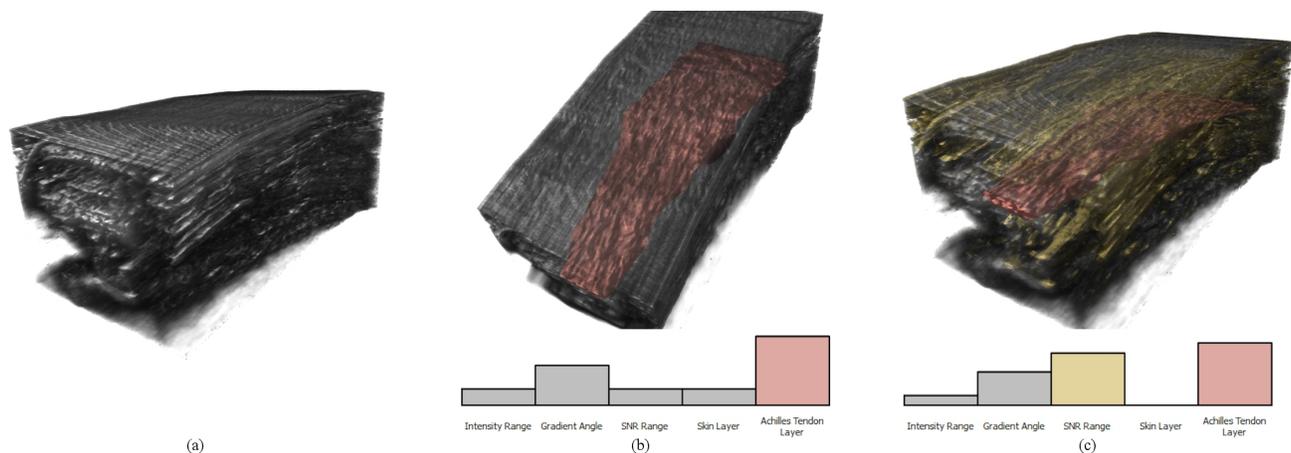


Fig. 8: Visualization results for the **achilles tendon data set** together with the used predicate histograms: Compared to the standard transfer function-based visualization (a), (b) highlights the target anatomy (achilles tendon) through the gradient angle predicate and the corresponding label predicate. (c) Shows the final rendering with highlighted fibrous muscle structures as context information using the SNR range predicate.

The resulting volumes have a resolution of 512^3 voxels for the shoulder data set and of 384^3 voxels for the carotid and the achilles tendon data set.

After discussions with clinicians we identified clinically relevant features for each data set: For the shoulder data set, the clinicians were interested in seeing the bone surface in context with the muscle layer. Figure 6 shows possible visualization results of our technique in conjunction with the predicate histogram used. 3D Visualization of carotid ultrasound specifically needs to show the path of the carotid artery and its bifurcation in a spatial context, as shown in Figures 1 and 7. Finally, for the achilles tendon acquisition, the clinicians need to see the tendon in its whole shape in order to identify possible tears or other lesions. Figure 8 shows the results of our approach.

To support our predicate-based approach to classification, we performed a manual segmentation on the data sets partitioning the ultrasound volumes into semantic layers of anatomy such as skin, muscle and bone (cf. Section 3.1 on label predicate).

To evaluate the transferability of the predicate histogram setup, we applied the same configuration to both our carotid ultrasound data sets. As depicted in Figure 9, a single predicate histogram was able to create a proper visualization for both data sets. This suggests that our formulation allows to reuse a predicate histogram, which was created for a specific data set, for a different data set of the same anatomy as valid preset, so that the user does not need to start the exploration and

tuning from scratch.

All results presented in this work were acquired in interactive sessions with a nVidia GeForce 670 GTX GPU. While implementing a performance-optimized rendering was not our main goal, our implementation is still capable of generating sufficient frame rates for interactive renderings as shown in Table 1. The decreased performance in comparison to standard 1D transfer function-based direct volume rendering is mainly because the illustrative focus-and-context rendering does no longer allow for early ray termination.

	Achilles Tendon (384^3 voxels)	Shoulder (512^3 voxels)
Standard 1D TF-based	102.3 fps	30.4 fps
Predicate-based Rendering	38.9 fps	15.7 fps

Table 1: Frame rates of our predicate-based approach compared to standard 1D transfer function-based direct volume rendering for different data sets on a GeForce 670 GTX, viewport size 800x600.

7 DISCUSSION

In the previous section, we compared our rendering results only to standard direct volume rendering with a 1D grayscale transfer function, which may raise the question how our technique compares to

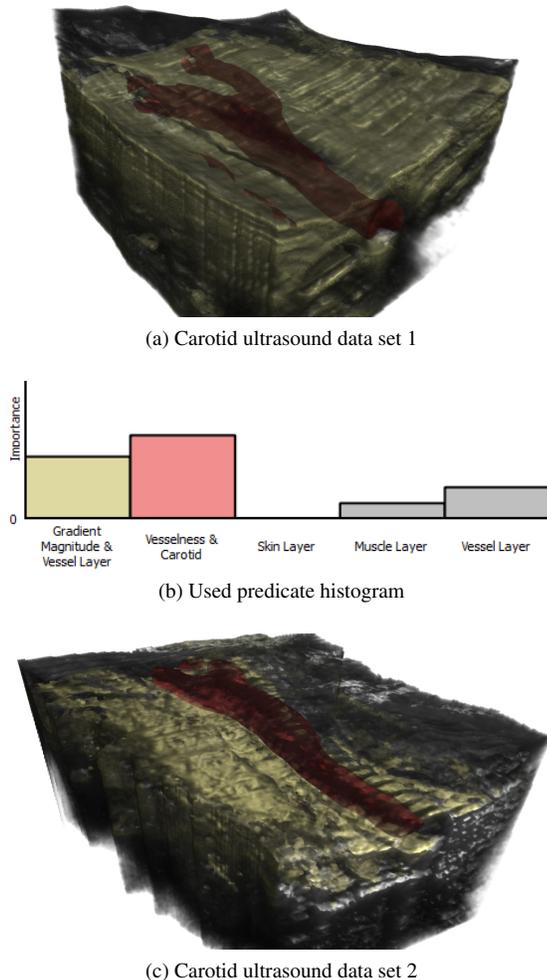


Fig. 9: Results suggest a **transferability of predicate histograms** between different data sets of the same anatomy: (a) A visualization for carotid data set 1 was created using the predicate histogram (b). This was then applied as preset to carotid data set 2 yielding a very viable visualization (c).

other related visualization techniques referenced in Section 2. The main issue is that most of these techniques are not applicable to ultrasound data. Ultrasound intensities are of highly context-sensitive nature (cf. Section 2.1) so that transfer functions solely relying on local information, such as intensity and gradient, neither improve the visual outcome nor are a suitable domain for defining color. Since our results rely to some extent on segmentation output, our approach can be seen as specialization of multi-volume rendering techniques, which could yield similar results. However, we would like to stress that the focus of our work is the full integration of a large variety of classification techniques into a single consistent formulation in combination with an intuitive user interface that allows for highly interactive exploration of the data.

One side effect of relevance-based rendering is the limited depth perception, in particular in still images. It is an inherent property of this approach that features, which would usually be occluded, pop out if they are assigned with a higher importance factor than features in the front. The correct depth perception can however still be perceived for instance through motion parallax in dynamic scenes as shown in our video in the supplementary material. Furthermore, the used compositing scheme [3] ensures a correct depth ordering in case of uniform importances so that the depth perception of the individual features themselves remains intact.

A very important aspect is the intuitiveness of our predicate his-

toqram based on the semantic nature of the predicates. Therefore, it is essential to provide the predicates and their settings with easily comprehensible naming. Since all the user needs to do is deciding whether he/she wants to see more or less of a certain feature and adjust the corresponding bar in the predicate histogram accordingly, this can also be considered as an implicit parameter space reduction

It should be noted that the focus of this paper is the introduction of the predicate concept and the predicate histogram, as well as the demonstration its potential on selected examples. An exhaustive study regarding the selection of suitable features for defining predicates is beyond the scope of this paper, as are techniques for automatic segmentation of specific anatomies. For a detailed overview on suitable segmentation techniques, we refer the interested reader to the book of Szeliski [24] or the survey paper of Heimann and Meinzer [7]. Due to our generic formulation, segmentation results obtained by such techniques or even the techniques themselves can be easily transferred into predicates and integrated into our visualization concept.

8 CONCLUSION

In this work, we introduced a novel consistent formulation for predicate-based classification of volumetric image data. It does not only allow for integration of both local and global image information and even anatomical models, but also naturally extends to illustrative focus-and-context visualization. While our technique does not make any a-priori assumptions on the type of input data or its spatial representation, it was designed for the specific application of direct volume visualization to 3D ultrasound volumes, for which it yields superior results in terms of occlusion and distinctly exposing selected image features than traditional global transfer function-based visualization.

To complement our approach, we additionally proposed the predicate histogram as an effective means for reducing the potentially high dimensionality of the predicate configuration domain.

It provides the user with an intuitive interface showing an overview over the parameter space, as well as with interaction metaphors to interactively manipulate the visualization result in real-time. Especially in combination with the implemented scribble technique, the predicate histogram can be modified by direct interaction with the rendered ultrasound image. Thus, we obtain an intuitive workflow, which allows also non-expert users to obtain insightful visualizations. The whole framework is designed to seamlessly integrate in the standard volume rendering pipeline without significantly increasing the computational burden and thus allowing for a real-time interaction with the system.

While directly integrating adequate segmentation technique into our predicate formulation might be one aspect of future work, other aspects are the application of our technique to other imaging modalities and multi-modal data sets. We imagine the predicate domain to be well suited for machine learning approaches, which could generate application-specific workflow models to automatically provide predicate histogram presets.

APPENDIX

The detailed recurrence scheme for the relevance-based compositing of our technique is as follows. The equations are the same as in [3] but adapted to our naming conventions.

Let $\kappa(s)$ be the computed sample importance, $C(s)$ be the sample color and $\alpha(s)$ be the sample opacity as we defined in Section 3.4. Then we define the visibility factor $\text{vis}(\kappa(s), \kappa_i)$ of the current sample $\kappa(s)$ compared to the accumulated importance κ_i as

$$\text{vis}(\kappa(s), \kappa_i) = 1 - \exp(-\kappa_i \kappa(s)), \quad (10)$$

as well as the modulation factor m as

$$m = \begin{cases} 1 & \kappa(s) \leq \kappa_i, \\ 1 - \alpha_i \geq \text{vis}(\kappa(s), \kappa_i), & \\ \frac{1 - \text{vis}(\kappa(s), \kappa_i)}{\alpha_i} & \text{else.} \end{cases} \quad (11)$$

The incremental front-to-back compositing scheme is defined by:

$$\begin{aligned} C'_{i+1} &= mC_i + (1 - m\alpha_i)C(s), \\ \alpha'_{i+1} &= m\alpha_i(1 - \alpha(s)) + \alpha(s), \\ \alpha_{i+1} &= \alpha_i(1 - \alpha(s)) + \alpha(s), \\ C_{i+1} &= \begin{cases} 0 & \alpha'_{i+1} = 0, \\ \frac{\alpha_{i+1}C'_{i+1}}{\alpha'_{i+1}} & \text{else,} \end{cases} \\ \kappa_{i+1} &= \max(\kappa_i, \ln(\alpha(s) + (1 - \alpha(s))\exp(\kappa_i - \kappa(s))) + \kappa(s)). \end{aligned} \quad (12)$$

ACKNOWLEDGEMENTS

This research was performed in the course of the *CAMPVis* project funded within the *SoftwareCampus* program of the German Federal Ministry of Education and Research (*BMBF, Förderkennzeichen 01IS12057*).

REFERENCES

- [1] S. Born, M. Pfeifle, M. Markl, M. Gutberlet, and G. Scheuermann. Visual Analysis of Cardiac 4D MRI Blood Flow Using Line Predicates. *Visualization and Computer Graphics, IEEE Transactions on*, 19(6):900–912, June 2013.
- [2] S. Bruckner, S. Grimm, A. Kanitsar, and M. E. Gröller. Illustrative Context-Preserving Volume Rendering. In K. Brodli, D. J. Duke, and K. I. Joy, editors, *EuroVis*, pages 69–76. Eurographics Association, 2005.
- [3] F. de Moura Pinto and C. Freitas. Importance-Aware Composition for Illustrative Volume Rendering. In *Graphics, Patterns and Images (SIBGRAPI), 2010 23rd SIBGRAPI Conference on*, pages 134–141, Aug 2010.
- [4] Q. Duan, E. D. Angelini, and A. F. Laine. Surface function actives. *Journal of Visual Communication and Image Representation*, 20(7):478–490, 2009.
- [5] R. Fattal and D. Lischinski. Variational classification for visualization of 3D ultrasound data. In *Visualization, 2001. VIS '01. Proceedings*, pages 403–410, Oct 2001.
- [6] A. Frangi, W. Niessen, K. Vincken, and M. Viergever. Multiscale vessel enhancement filtering. In W. Wells, A. Colchester, and S. Delp, editors, *Medical Image Computing and Computer-Assisted Intervention - MICCAI*, volume 1496 of *Lecture Notes in Computer Science*, pages 130–137. Springer Berlin Heidelberg, 1998.
- [7] T. Heimann and H.-P. Meinzer. Statistical shape models for 3D medical image segmentation: A review. *Medical Image Analysis*, 13(4):543–563, 2009.
- [8] C. Hennersperger, D. Mateus, M. Baust, and N. Navab. A Quadratic Energy Minimization Framework for Signal Loss Estimation from Arbitrarily Sampled Ultrasound Data. In *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2014*, Boston, USA, September 2014.
- [9] A. Karamalis, W. Wein, T. Klein, and N. Navab. Ultrasound confidence maps using random walks. *Medical Image Analysis*, 16(6):1101–1112, 2012.
- [10] A. Katouzian, E. Angelini, B. Sturm, and A. Laine. Brushlet segmentation for automatic detection of lumen borders in ivus images: A comparison study. In *Biomedical Imaging (ISBI), 2012 9th IEEE International Symposium on*, pages 242–245, May 2012.
- [11] J. Kniss, G. Kindlmann, and C. Hansen. Multidimensional transfer functions for interactive volume rendering. *Visualization and Computer Graphics, IEEE Transactions on*, 8(3):270–285, Jul 2002.
- [12] B. Kohler, R. Gasteiger, U. Preim, H. Theisel, M. Gutberlet, and B. Preim. Semi-Automatic Vortex Extraction in 4D PC-MRI Cardiac Blood Flow Data using Line Predicates. *Visualization and Computer Graphics, IEEE Transactions on*, 19(12):2773–2782, Dec 2013.
- [13] J. Kruger, J. Schneider, and R. Westermann. ClearView: An Interactive Context Preserving Hotspot Visualization Technique. *Visualization and Computer Graphics, IEEE Transactions on*, 12(5):941–948, Sept 2006.
- [14] D. Mann, J. J. Caban, P. J. Stolka, E. M. Boctor, and T. S. Yoo. Multidimensional TF for effective visualization of streaming US and elasticity images. volume 7964, pages 796439–796439–10, 2011.
- [15] J.-S. Prassni, T. Ropinski, J. Mensmann, and K. Hinrichs. Shape-based transfer functions for volume visualization. In *Pacific Visualization Symposium (PacificVis), 2010 IEEE*, pages 9–16, March 2010.
- [16] B. Preim and C. Botha. *Visual Computing for Medicine: Theory, Algorithms, and Applications*. The Morgan Kaufmann Series in Computer Graphics. Elsevier Science, 2013.
- [17] P. Rautek, S. Bruckner, and M. E. Gröller. Interaction-dependent Semantics for Illustrative Volume Rendering. In *Proceedings of the 10th Joint Eurographics / IEEE - VGTC Conference on Visualization, EuroVis'08*, pages 847–854, Aire-la-Ville, Switzerland, Switzerland, 2008. Eurographics Association.
- [18] T. Ropinski, J. Prassni, F. Steinicke, and K. Hinrichs. Stroke-based Transfer Function Design. In *Proceedings of the Fifth Eurographics / IEEE VGTC Conference on Point-Based Graphics, SPBG'08*, pages 41–48, Aire-la-Ville, Switzerland, Switzerland, 2008. Eurographics Association.
- [19] G. Sakas, L.-A. Schreyer, and M. Grimm. Preprocessing and volume rendering of 3D ultrasonic data. *Computer Graphics and Applications, IEEE*, 15(4):47–54, Jul 1995.
- [20] C. Salama, M. Keller, and P. Kohlmann. High-Level User Interfaces for Transfer Function Design with Semantics. *Visualization and Computer Graphics, IEEE Transactions on*, 12(5):1021–1028, Sept 2006.
- [21] T. Salzbrunn, C. Garth, G. Scheuermann, and J. Meyer. Pathline predicates and unsteady flow structures. *The Visual Computer*, 24(12):1039–1051, 2008.
- [22] C. Schulte zu Berge, A. Kapoor, and N. Navab. Orientation-Driven Ultrasound Compounding Using Uncertainty Information. In *IPCAI*, pages 236–245, 2014.
- [23] P. Sereda, A. Bartroli, I. W. O. Serlie, and F. Gerritsen. Visualization of boundaries in volumetric data sets using LH histograms. *Visualization and Computer Graphics, IEEE Transactions on*, 12(2):208–218, March 2006.
- [24] R. Szeliski. *Computer Vision: Algorithms and Applications*. Springer-Verlag New York, Inc., New York, NY, USA, 1st edition, 2010.
- [25] F.-Y. Tzeng, E. B. Lum, and K.-L. Ma. A Novel Interface for Higher-Dimensional Classification of Volume Data. In *Proceedings of the 14th IEEE Visualization 2003 (VIS'03)*, VIS '03, pages 66–, Washington, DC, USA, 2003. IEEE Computer Society.
- [26] Y. Wan and C. Hansen. Fast Volumetric Data Exploration with Importance-based Accumulated Transparency Modulation. In *Proceedings of the 8th IEEE/EG International Conference on Volume Graphics, VG'10*, pages 61–68, Aire-la-Ville, Switzerland, Switzerland, 2010. Eurographics Association.
- [27] L. Wang, X. Zhao, and A. Kaufman. Modified Dendrogram of Attribute Space for Multidimensional Transfer Function Design. *Visualization and Computer Graphics, IEEE Transactions on*, 18(1):121–131, Jan 2012.