

Uncertainty Visualization in Medical Volume Rendering Using Probabilistic Animation

Claes Lundström, *Student Member, IEEE*, Patric Ljung, *Member, IEEE*, Anders Persson, and Anders Ynnerman, *Member, IEEE*

Abstract— Direct Volume Rendering has proved to be an effective visualization method for medical data sets and has reached wide-spread clinical use. The diagnostic exploration, in essence, corresponds to a tissue classification task, which is often complex and time-consuming. Moreover, a major problem is the lack of information on the uncertainty of the classification, which can have dramatic consequences for the diagnosis. In this paper this problem is addressed by proposing animation methods to convey uncertainty in the rendering. The foundation is a probabilistic Transfer Function model which allows for direct user interaction with the classification. The rendering is animated by sampling the probability domain over time, which results in varying appearance for uncertain regions. A particularly promising application of this technique is a “sensitivity lens” applied to focus regions in the data set. The methods have been evaluated by radiologists in a study simulating the clinical task of stenosis assessment, in which the animation technique is shown to outperform traditional rendering in terms of assessment accuracy.

Index Terms—Uncertainty, probability, medical visualization, volume rendering, transfer function.

1 Introduction

Modern visualization techniques have successfully been applied to diagnostic medical imaging and Direct Volume Rendering (DVR) has in recent years matured into a routinely used tool in the clinical work. Nevertheless, the diagnostic work still presents significant challenges where existing DVR approaches are insufficient. Uncertainty visualization is one such challenge and it is the focus of this paper.

The simplest way of working with DVR is to use static, pre-defined Transfer Functions (TFs), referred to as *presets*. This is common in actual clinical usage for CT data sets, since the sample range is considered to be calibrated (the Hounsfield scale). Experienced radiologists have learned, however, to always question the presets, even in the CT case. The reason is that individual variations between patients and diagnostic tasks often make pre-defined TFs inadequate. One example is the false stenosis in figure 1, where a patient underwent an unnecessary surgical procedure due to a misleading preset TF. An example of an even more serious situation is pre-operative assessment of Abdominal Aortic Aneurysms. If the vessel width is incorrectly determined from the image, a stent of the wrong size may be put into the aorta which is directly life-threatening. The fact that inadequate visualization parameters can change the apparent vessel width dramatically [15] underlines the need of exploring all feasible visualization settings.

Therefore, radiologists employ repeated manual TF adjustments to study several possible visualization outcomes to be considered in the diagnosis. This routine is an ad hoc uncertainty assessment of the visualization parameters. A major drawback is that the manual adjustments are very time-consuming. Another important issue is that the manual adjustments are not performed in a controlled manner. Even with moderately complex TFs, it is unlikely that the radiologist’s interaction explores the full parameter space relevant to the diagnosis.

There is a great risk that the process becomes far from the ideal unbiased exploration of all possibilities, in particular for physicians with limited experience of DVR.

In this paper uncertainty visualization methods are proposed that aim to increase the accuracy and efficiency of DVR as a diagnostic tool. Uncertainty visualization often deals with pre-classified data sets, i.e. the probability values exist before the visualization process starts. As the clinical routines show, however, it is highly relevant to explore the uncertainty inherent in the TF classification. From this perspective, it is necessary to distinguish between the two roles of the TF: the tissue classification and the rendering appearance control. An explicitly probabilistic TF model is presented in this work, which separates the two roles and provides a direct control of the classifying criteria.

Based on the probabilistic TF formulation, a number of animation schemes to convey relevant uncertainty are defined. Image regions corresponding to confident tissue classifications will be static, whereas uncertain parts will change over time. For each sample, the probability of a material is reflected by the number of frames in which it appears in the animation cycle.

This paper contributes to several aspects of uncertainty visualization research. A novel method to integrate animation in DVR is presented. This solution avoids drawbacks of previous methods targeting uncertainty, such as being limited to data with distinct surfaces or suffering from decreased spatial precision when conveying indecisive regions. Furthermore, by isolating the probabilistic part of the TF, the fuzziness of the material classification is given a form that enables advanced visualization and interaction. Finally, the techniques presented constitute valuable tools in clinical use of DVR. This is demonstrated by a user study showing that the method outperforms traditional rendering in a simulated stenosis assessment task.

The problem description is based on discussions with radiologists routinely using DVR at the Center for Medical Image Science and Visualization (CMIV), Linköping University. The patients referred to CMIV are typically complex cases, requiring elaborate diagnostic assessment using state-of-the-art Computed Tomography (CT) and Magnetic Resonance (MR) imaging methods. In many of these cases, DVR is used as a diagnostic tool.

2 Related work

Uncertainty visualization is often brought forward as one of the main challenges for visualization research [9]. In this section the previous work most relevant to this paper is reviewed. For a broader discussion on the cognitive theories concerning uncertainty visualization, refer to the recent survey by Zuk and Carpendale [18].

- Claes Lundström is with the Center for Medical Image Science and Visualization (CMIV), Linköping University, and Sectra-Imtec AB. E-mail: clalu@cmiv.liu.se.
- Patric Ljung and Anders Ynnerman are with the Division for Visual Information Technology and Applications (VITA), Linköping University. E-mail: {plg, andyn}@itn.liu.se.
- Anders Persson is with the Center for Medical Image Science and Visualization (CMIV), Linköping University. E-mail: anders.persson@cmiv.liu.se.

Manuscript received 31 March 2007; accepted 1 August 2007; posted online 27 October 2007. Published 14 September 2007.

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

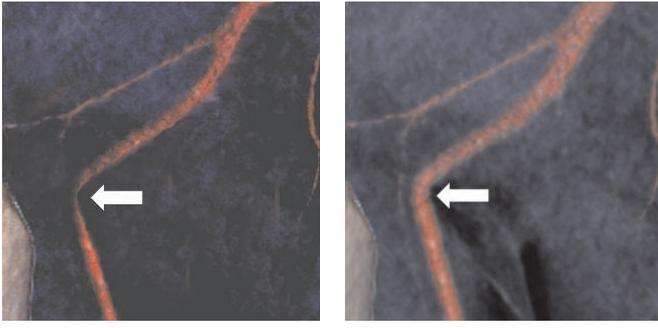


Fig. 1. Static TF presets can lead to incorrect conclusions, even for calibrated intensities as in this CT angiography. The renderings were created by commercial DVR software in clinical use. Left: The rendering using the standard TF indicates a stenosis, pointed out by the arrow. Based on this image, an unnecessary surgical procedure was performed. Right: Exploring other TF settings clearly shows that there is no stenosis.

There are a number of examples of uncertainty integrated in surface renderings. Grigoryan and Rheingans [6] employed a point cloud approach, where points were offset from the estimated surface proportionally to the level of uncertainty. The work of Kindlmann et al. [10] includes the visualization of shape uncertainty of an isosurface through flowline curvature. Rhodes et al. [17] targeted the uncertainty introduced by multiresolution models, mapping the isosurface confidence to hue or texture. Finally, a broad range of uncertainty representations for surface renderings was presented by Pang et al. [14].

The above methods have limited use for volumetric data sets where there are no distinct surfaces and so methods tailored for DVR are needed. Rheingans and Joshi [16] addressed positional uncertainty in molecule models, among other methods proposing rendering of the likelihood domain itself. Djurcilov et al. [1] presented several approaches: probability values were mapped to opacity or used as an additional variable in a 2D TF. Moreover, a separate rendering of the probability volume was used to modulate the pixels of the rendered image. A fundamental limitation of the latter approach is that regions with high uncertainty can affect the rendering even if they are obscured in the regular visualization. Kniss et al. [11] proposed a framework for DVR of volumes with predefined probabilistic classification. The approach is based on statistical risk and both DVR and surface rendering applications were shown.

Turning to the classification aspect of the TF, the currently dominating TF approach is arguably to *implicitly* model material probability. TF models where material probabilities are *explicitly* derived are more rare, even though DVR was first developed with such a model by Drebin et al. [2]. The few examples of recent DVR methods explicitly dealing with the probability domain include the work of Lundström et al. [12, 13], where a probabilistic model for pair-wise material separation was proposed that can be used to convey uncertainty. Furthermore, the visualization techniques of Kniss et al. [11] convey quantitative statistical information and operate in the material probability domain, but the methods are restricted to pre-classified data sets.

Motion perception is a prioritized preattentive process of human vision. Animation schemes for uncertainty visualization have been proposed in several geographical applications. Gershon [5] proposed to animate an ordered set of segmentations to enhance visibility of fuzzy structures. Another example is the work of Ehlschlaeger et al. [3], where spatial uncertainty was conveyed as a sequence of probabilistic realizations.

Many types of uncertainty sources have been considered in previous work. In many cases, the probability values are assumed to exist before the visualization process starts. In contrast, the focus of this paper is uncertainty arising from the TFs fuzzy classification within DVR.

3 Probabilistic Uncertainty Animation

An initial issue in uncertainty visualization applications is to define a relevant type of uncertainty. In the medical context, the DVR application is traditionally provided with an unclassified data set. The experience of the radiologists' clinical DVR usage shows that the TF is in practise used as an interactive tool for fuzzy classification of tissues. The methods proposed in this paper aim to formalize and automate the visualization of these classification uncertainties.

The classifying role of the TF defines the information needed for the uncertainty analysis. Therefore, this part needs to be separated from the second role, the appearance mapping. To this end, a probabilistic TF model is defined in section 3.1. The model is then employed in a method for probabilistic animation, described in section 3.2. The fundamental idea is that a sequence of renderings is an adequate representation of information that cannot be conveyed in any single, static image. Three different modes of probabilistic animation are presented in section 3.3. Finally, an application of these techniques tailored for the clinical work flow, the *sensitivity lens*, is described in section 3.4.

At a general level, the benefits of the methods proposed in this work can be summarized as follows. The uncertainty visualization is integrated in the DVR, as opposed to being an overlay on the standard rendering. As each frame is a familiar DVR image, the uncertainty information can be easily grasped by the user. The animation approach provides a fast way of studying many alternative renderings. Another important property is that spatial precision is not decreased.

3.1 Explicitly Probabilistic Transfer Function

In the context of this paper, two main views on TFs can be distinguished. The TF is seen as a classification tool in both views, 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 $\mathbf{c}(s)$. Such an approach was employed in the initial DVR implementation by Drebin et al. [2], according to equations 1 and 2. Transparent regions are referred to as the *null material* ($m = 0$).

$$\mathbf{c}(s) = \sum_{m=0}^M p_m(s) \cdot \tilde{\mathbf{c}}_m \quad (1)$$

$$\tilde{\mathbf{c}}_m = (\alpha_m r_m, \alpha_m g_m, \alpha_m b_m, \alpha_m)^T \quad (2)$$

In the implicitly probabilistic view the TF is seen as a direct mapping from sample value s to sample color $\mathbf{c}(s)$. This is the currently dominating approach and it is the view represented in recent DVR literature [7, 4]. Fuzzy classification is typically achieved by connecting material probability to the opacity level, giving low opacity to uncertain materials. In the terms of equation 1, the product $p_m(s) \cdot \alpha_m$ is integrated into $\alpha'_m(s)$.

The objective of the uncertainty visualization in this work is to explore relevant alternative renderings, given a TF. The implicit representation is not suitable for this task. A crude form of sensitivity analysis can be made by perturbing the TF parameters, but there is no theoretical base providing control of the exploration. In contrast, the explicit probabilistic representation provides a well-defined uncertainty domain.

Therefore, the proposed TF model is explicitly probabilistic, defined as follows. Each material to be visualized is connected to an individual TF component consisting of two parts, as shown in figure 2. The first is the material appearance ($\text{rgb}\alpha$), which in this work is chosen to be static in order to promote simplicity in the interpretation of the rendered image. The second part is the classifying function $\tilde{p}_m(s)$ that maps intensity to material likelihood. The intensity-specific material probabilities are the normalized likelihoods:

$$p_m(s) = \frac{\tilde{p}_m(s)}{\sum_{m'=0}^M \tilde{p}_{m'}(s)} \quad (3)$$

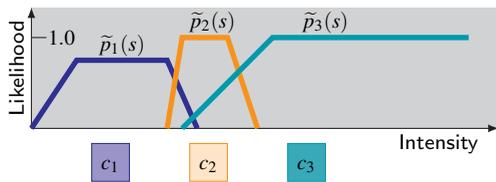


Fig. 2. In the explicitly probabilistic TF model, the familiar TF GUI components define the material likelihood. The material appearance is separated from the classification process.

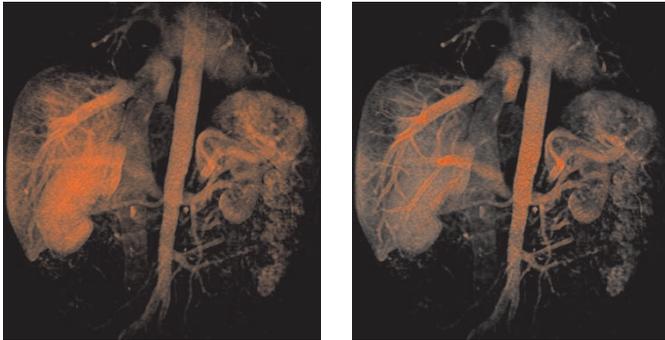


Fig. 3. Modeling probability by desaturation for an MR renal angiography. Left: Traditional rendering, where it is hard to distinguish between thin tissue regions and regions with low tissue classification probability. Right: Desaturating uncertain regions using a probabilistic DVR approach. The TF consists of a single component describing the likelihood of vessel tissue.

The user defines $\tilde{p}_m(s)$ in a TF GUI, a convenient form is trapezoids in the $[0.0, 1.0]$ likelihood range. Null material likelihood is implicitly defined:

$$\tilde{p}_0(s) = \max(0.0, 1.0 - \sum_{m=1}^M \tilde{p}_m(s)) \quad (4)$$

In some cases it is useful to compare standard and probabilistic TFs. Therefore, a mapping between the two types has been developed for TFs with material-specific components. The proposed mapping is not guaranteed to give useful results in all cases, but it is feasible for certain TF types. The material appearance \mathbf{c}_m is set to $(r_m, g_m, b_m, \hat{\alpha}_m)^T$, where $\hat{\alpha}_m$ is the maximum opacity of the material. The material likelihood is a simple scaling of the opacity profile:

$$\tilde{p}_m(s) = \alpha'_m(s) / \hat{\alpha}_m \quad (5)$$

The benefits of the explicitly probabilistic TF model extend beyond being the foundation of the proposed uncertainty animation, described in section 3.2. Having a separate probability component enables many possibilities for visualization of statistical properties, also for non-animated renderings. An example is to connect uncertainty to color desaturation, which has previously been proposed in other domains [8]. With explicit probabilities, it is straightforward to achieve this effect, see figure 3. Furthermore, working in a material likelihood GUI can benefit the physicians' understanding of the TF adjustment process as an exploration in form of classification interaction rather than as an appearance modification.

3.2 Probabilistic Animation

In this section animation techniques for exploring the domain of possible outcomes from a probabilistic classification are proposed. In essence, the approach translates to a sequence of Monte Carlo simulations of the rendering. The classifier in focus is the probabilistic TF model presented in section 3.1, resulting in a set of material probability values $p_0(s), \dots, p_M(s)$ for each intensity value s . The animation

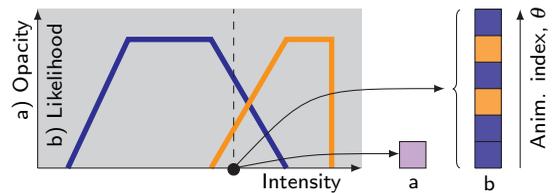


Fig. 4. Traditional DVR vs. probabilistic animation. a) Traditionally, the TF maps an intensity value to a single color, a mixture of material colors. b) The probabilistic animation maps an intensity value to an array of pure material colors.

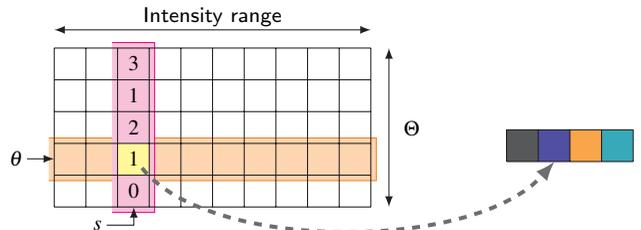


Fig. 5. Fragment shader implementation of probabilistic uncertainty animation. From left to right: The sample intensity value s is a column index into the material selection LUT. Together with the animation index θ , a material index is looked up and finally used to find the resulting sample color.

techniques can, however, be employed in combination with an arbitrary probabilistic classification method.

The key part of the animated rendering is the derivation of the sample color $\mathbf{c}(s)$ from the material probabilities. In a traditional rendering, the process would be to mix the materials' colors, weighted by the probabilities as described by equation 1. Introducing animation, the possibility to represent probability in the time dimension is added. A straight-forward linear mapping is employed: Having $p_m(s) = x\%$ translates to setting $\mathbf{c}(s) = \mathbf{c}_m$ in $x\%$ of the frames in the animation. Doing this for all materials, the sequence of sample colors captures the probabilistic properties of the sample. Applying this for all samples in the volume will, consequently, capture the probabilistic properties of the entire data set, see figure 6. Intensity values corresponding to uncertain material classification will have varying color mapping in the rendering. Note that to achieve a relevant probabilistic representation in the animation, the material colors cannot be mixed in the sampling stage. Color mixtures in the rendering only arise from depth-wise compositing.

Let Θ denote the total number of steps in the animation cycle. Thus, the probabilistic animation scheme should map each intensity value to an array of Θ material colors, in contrast to the single $\text{rgb}\alpha$ value in the traditional case. This is illustrated in figure 4. Expanding this to the whole intensity range, the normal 1D look-up of material color mixtures is replaced by a 2D look-up of pure material colors. The added dimension is named the *animation index*, θ . The animation is achieved by cyclically changing the θ value for each rendered frame. The resolution in the probability domain is controlled by the Θ value. The renderings of medical data sets in this paper have $\Theta = 16$. The number of slots for each material in the animation cycle, $n_m(s)$, see eq. 6, must be transformed into an integer value. This is done in an iterative approach: assign a slot to the material having most slots, \hat{m} , see eq. 7, then decrease $n_{\hat{m}}(s)$ by 1.0 and restart until all slots are filled.

$$n_m(s) = p_m(s) \cdot \Theta \quad (6)$$

$$\hat{m} = \operatorname{argmax}_m(n_m(s)), m = \{0, 1, \dots, M\} \quad (7)$$

The animation has been implemented as part of a GPU-based texture-slicing volume renderer. The only modification of the standard

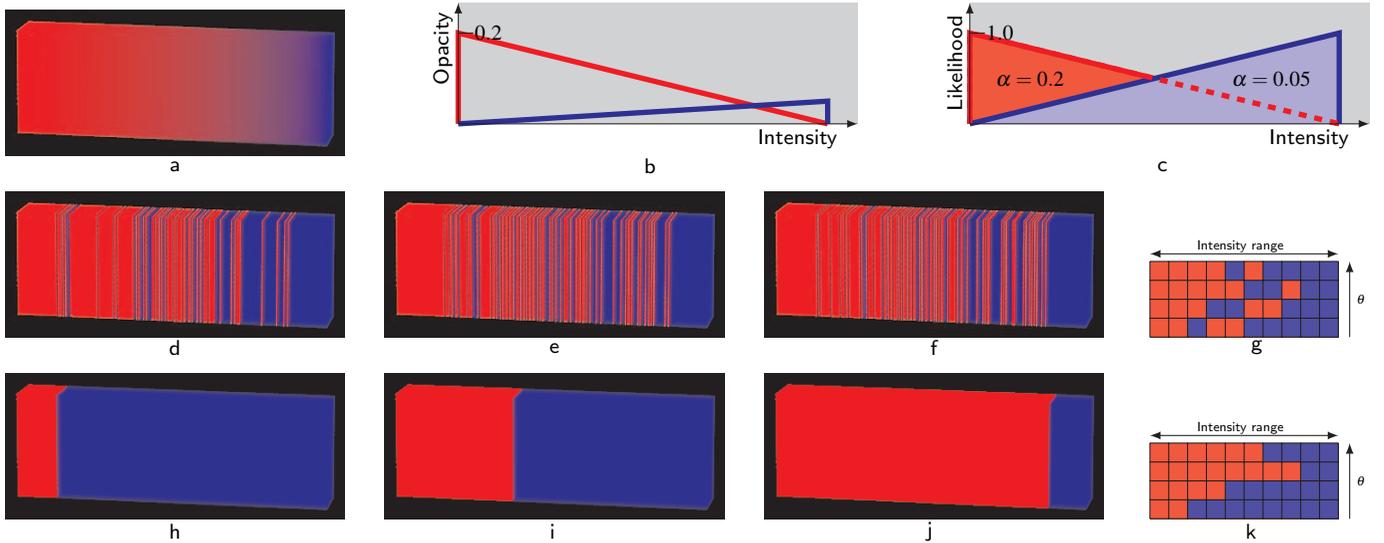


Fig. 6. Probabilistic animation. The data set consists of a linear gradient in the horizontal direction. a) Traditional rendering with a clear bias towards the high-opacity red material. b) Corresponding traditional TF. c) Explicitly probabilistic TF used for the animations. d-f) Frames from probabilistic animation ($\Theta = 4$) in random mode, for $\theta = 0, 1, \text{ and } 2$. g) Schematic view of intensity \Leftrightarrow color mapping in random mode. h-j) Frames from probabilistic animation ($\Theta = 4$) in material sync mode, for $\theta = 0, 1, \text{ and } 2$. k) Schematic view of intensity \Leftrightarrow color mapping in material sync mode.

implementation is to replace the color-opacity look-up table (LUT) by a material selection LUT and a material color array, see figure 5. The material selection LUT is a 2D table, where each row spans the intensity range and the columns correspond to the Θ steps in the animation cycle. The table is filled with material indices. Thus, each row represents a specific intensity \Leftrightarrow material mapping for an animation frame. The found material index points to the appropriate position of the material color array, which yields the fragment color.

The material selection LUT is derived in software as soon as the TF is changed. For each frame to render the animation index θ is updated, which is used to sample the correct row of the material selection LUT.

The presented methods affect all samples of a certain intensity equally. An alternative would be to perform an individual material color lookup for each sample, i.e. changing θ between samples. This entails, however, disadvantages in terms of computational complexity and reduced difference between animation frames, which is why this path has not been pursued.

This paper focuses on the simple classification of a 1D intensity TF. The general case is that an arbitrary classification scheme has produced a volume with material probabilities. The probabilistic animation is expected to generalize well, a theoretical discussion is given in appendix A.

3.3 Animation Modes

As described above, the material probabilities connected to a certain intensity value are represented by the relative number of entries of each material in the corresponding column of the material selection LUT. The term “animation array” will be used for such a column. The probabilistic relevance is retained as long as the number of entries for each material is not changed, but the *order* of the entries can be changed without losing the probability mapping. Different uncertainty visualization effects can be achieved by elaborate arrangement within the array. Three animation modes have been defined (see figures 6 and 7):

- **Random mode.** The population of the array is fully randomized.
- **Grouped random mode.** The entries are grouped for each material, but the order and position of the groups are randomized.
- **Material sync mode.** Each material has a fix base position in the array, around which the material entries are centered.

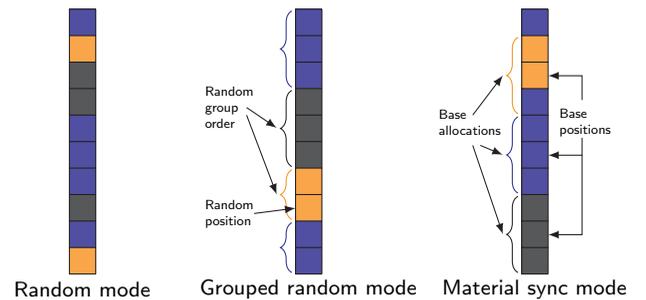


Fig. 7. The three proposed probabilistic animation modes. The modes are defined by the arrangement of material entries across the animation index dimension, here illustrated for a single intensity value. Note that the number of entries for each material is the same across all modes. (Black = null material) Left: Random mode. The materials are randomly inserted according to the probabilities. Middle: Grouped random mode. The entries are grouped according to the material and the groups are randomly placed. Right: Material sync mode. The materials are centered at a respective base position.

The random mode is a straight-forward approach achieving the probabilistic animation effect. The motivation for the grouped random mode is that it has a smoother appearance than the random mode, which mitigates the visual fatigue of looking at a flickering image. The benefit of the material sync mode is that there will be an animation frame that shows the maximum extent of each material, at the animation index corresponding to the material’s base position. This synchronization requires a sorting procedure for each value in the sample range and is performed as follows. The algorithm should produce an array containing Θ material selections for each intensity value. A common base allocation span d and material base positions u_m are set:

$$d = \lfloor \Theta / (M + 1) \rfloor, \quad u_m = m \cdot d + \lfloor \frac{d}{2} \rfloor, \quad (8)$$

Materials are entered into the array at the position closest to the base position. The base allocation is reserved for the corresponding material for its first d selections, if there are so many. When all possible base allocations are done, the remaining empty slots are filled with

materials having a number of entries exceeding the base allocation.

Variations of the animation techniques can be put to use in different scenarios. An immediate, qualitative impression of the uncertainty is achieved by the random or grouped random modes at a fairly high animation speed (e.g., 10 fps). For more detailed analysis of possible extents of a tissue, the material sync mode at a moderate animation speed (e.g., 5 fps) is appropriate. Another possibility is to use the material sync mode for TF tuning. Using a low animation speed (e.g., 2 fps), the user can fix the TF at the animation step that best represents the features of interest.

3.4 Sensitivity Lens

A crucial part of the development of visualization schemes is to tailor the general techniques to fit the specific application context. In the clinical case, the usage needs to meet extreme demands on simplicity and efficiency since the physicians have very limited time for each diagnosis. A fact to be exploited for uncertainty animation is that this advanced exploration is seldom needed for the entire volume. The typical case is rather that a sensitivity analysis is wanted for a few essential regions.

Based on this observation, the *sensitivity lens* application is proposed, see figure 10. The lens consists of a small user-controlled region-of-interest in which the selected animation technique is applied while the rest of the image is rendered traditionally. This tool is expected to be easy to master, since it can be used without parameter tuning and since it resembles the familiar magnifying glass tool. Another advantage is that potential visual fatigue from studying the animations is reduced thanks to the small region.

The sensitivity lens has potential to become an important tool in the clinical DVR work flow. As a consequence of the high radiologist workload, it is not uncommon that the technician scanning the patient also works as the 3D expert. The technician sets the visualization parameters and the radiologist merely rotates the volume during the diagnostic review. There are significant concerns for diagnosis quality with this approach, since the technician's mistakes may directly influence the radiologist's assessment. The sensitivity lens can in this situation act as a quality assurance tool, enabling the radiologist to explore alternative renderings without requiring additional usage complexity in terms of TF modification.

4 Results

The tests of the proposed animation techniques evaluate aspects of the clinical usefulness. The results are divided into two parts. First, results from an experimental study on stenosis assessment for simulated vessels are shown. The second part consist of renderings of data sets from diagnostic situations where the classification task of DVR is particularly challenging and sensitive to errors. Note that the still image presentation in this paper makes the frame-to-frame differences appear overly subtle. Refer to the video file available at <http://www.cmiv.liu.se/Members/clalu/uncert.mpg> for a correct impression of the animation.

4.1 Simulated Clinical Task

A clinically relevant test of the benefits of uncertainty animation is to find out whether it can increase the diagnostic accuracy in the challenging case of MR angiographies. Therefore, an experimental study on simulated vessels was carried out, see figure 8. The test subjects were twelve physicians, eleven radiologists and one cardiologist, all with clinical experience of stenosis assessment.

The test was designed to resemble a low-quality data set in combination with an untuned and fixed TF, when zooming in on a suspected stenosis. Three rendering methods were tested for 24 vessels: traditional static rendering and uncertainty animation in grouped random and material sync modes. (To minimize the effort of the test subjects, the random mode was left out since it was assessed to be very similar to grouped random mode in this test setup.) The animations were fixed at 6.0 fps. The accuracy of each method was measured as the absolute error in assessing the stenoses on a four-grade scale. The details of the study setup is given in appendix B.

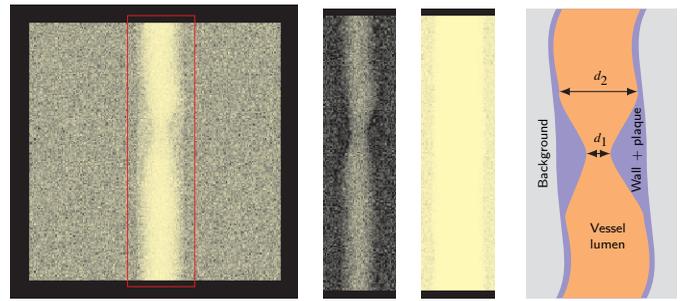


Fig. 8. Example of simulated stenosis assessment. The outer perimeter of the vessel model has constant radius, whereas there is a concentric plaque that creates a stenosis. The plaque profile is a cosine curve. The task is to assess the stenosis degree, d_1/d_2 . The three materials vessel lumen, vessel wall/plaque, and background have overlapping intensity distributions in order to make the task challenging. From left to right: An example of a rendered frame in material sync mode. Cut-outs (red frame of full image) of vessel at two other time steps. A schematic model.

Table 1. Error Scores and Comparisons for Stenosis Assessment

Method	Median	IQR
Material Sync	7.5	1.0
Grouped Random	13	3.5
Traditional	14	7.5

Comp.	p
Material Sync < Traditional	0.0057
Material Sync < Grouped Random	0.013
Grouped Random < Traditional	0.35

The result of the study is presented in table 1, method results at the top and pairwise comparisons at the bottom. The scores are aggregated assessment error, a lower score means a higher accuracy. The score spread is presented as interquartile range (IQR). The statistical significance level was set at $p < 0.05$ and the statistical analysis was performed in StatView 5.0 (SAS Institute, Cary NC, USA).

The result shows that the material sync mode provides the highest accuracy in this task. The aggregated error for each subject/method combination was analyzed using Friedman's test, which showed that there is significant differences between the methods ($p = 0.0069$). Subsequent pairwise Wilcoxon tests showed that the material sync mode has significantly lower error (higher accuracy), than the other two methods. The grouped random animation has slightly lower median than the traditional rendering, but the difference is not significant.

In order to establish a "gold standard" benchmark, the subjects also assessed the 24 vessels in a traditional rendering with free manual adjustment of the TF. The median accuracy was found to be 6, which means that the material sync mode, with a median of 7.5, comes quite close to the gold standard even though no TF adjustment was allowed in that case.

A brief interview was carried out with the physicians after the test. Out of the tested methods, ten out of the twelve preferred the material sync mode in the test setup. Two reasons repeatedly came up as motivation: that the vessel structure was clearly seen and that the synchronization resembled the familiar manual back-and-forth TF adjustment. Two physicians preferred the grouped random mode, but their accuracy was actually lowest for that method. Furthermore, there was consensus that uncertainty animation would improve accuracy and/or efficiency for difficult stenosis assessments. Finally, nearly all physicians stressed that in real clinical usage the physician must be able to interactively control the uncertainty animation in terms of the speed and the underlying probabilistic TF.

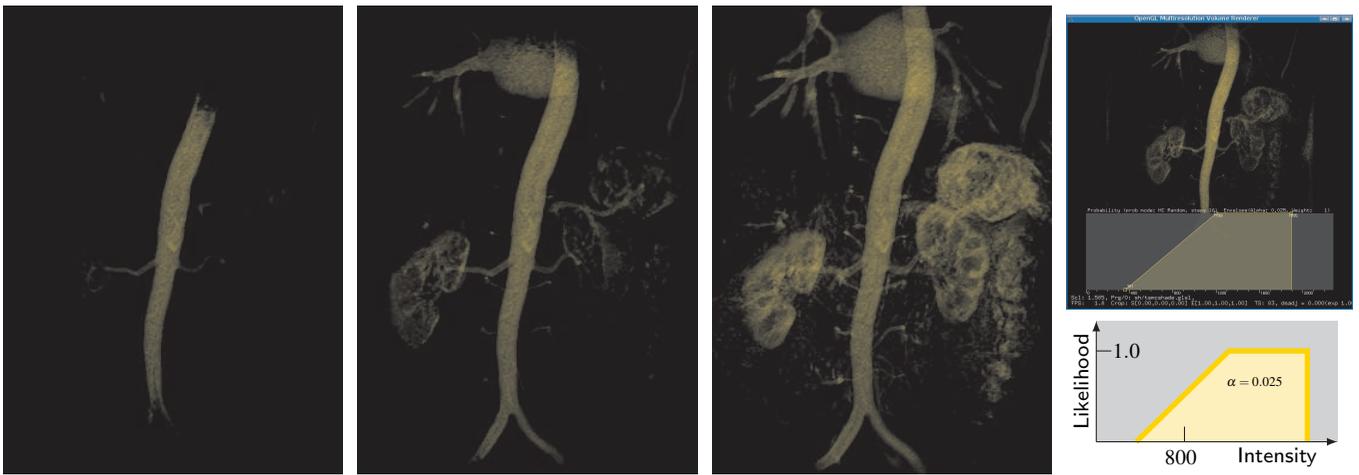


Fig. 9. Uncertainty animation in material sync mode for a MR renal angiography. Three examples of animation frames connecting visibility to different vessel probability levels. Rightmost, top: A screen shot from the application GUI. Bottom: The probabilistic TF.

4.2 Renderings of Clinical Data Sets

Radiologists consider MR angiographies to be a clinical examination with great need for uncertainty visualization. MR examinations are preferred over CT scans since ionizing radiation is potentially damaging and is to be avoided if possible. DVR of MR angiographies is, however, time-consuming and error-prone. Static TFs can not be used as the intensity scale varies between patients and an inadequate visualization can give an incorrect impression of the vessels that affects the medical assessment [15]. The animated renderings have been applied to several MR angiographies, two examples are shown next.

Figure 9 shows the uncertainty animation in material sync mode for a MR renal angiography. The diagnostic task is to determine location and significance of any stenoses. The vessels cannot be isolated through any single setting of an intensity-based TF; when the minor vessels appear, so do the kidneys and other parts. With the uncertainty animation, the alternative visualizations can be explored by a single TF setting.

The *sensitivity lens* was of particular interest for clinical use. A stenosis assessment application for another MR renal angiography is shown in figure 10. If a non-blocked vessel is shown in one of the alternative frames, this means that there is no stenosis, merely a region of lower contrast agent concentration. Many radiologists in the user study of section 4.1 commented that uncertainty animation would be most useful as an additional check for regions of interest. In such a focused review, the sensitivity lens was judged to be a highly suitable tool.

Despite the calibrated Hounsfield scale, CT data sets often have overlapping tissue intensity ranges which can cause ambiguous visualizations. Such an example is the thyroid tumor in figure 11, where the tumor extent in relation to the left carotid arteries is crucial to the pre-operative planning. The highly different frames show that there may be a large span of alternative visualizations even though the traditional, static rendering seems distinct.

The final example is an MR brain examination showing a large cyst, see figure 12. The important clinical question is if the cyst is distinctly separated from the hemorrhage, the two regions having very similar intensities. The uncertainty animation provides a controlled exploration of the extent of the tissues.

On the tested system, the rendering performance was about 10% lower for the animation methods compared to the standard method, see table 2. The derivation of the material selection LUT required 1-3 ms. The test system was a PC with an AMD Athlon 64 CPU and an nVidia 8800 GTX graphics board with 768 MB, the volumes were rendered in a 512x512 view port.

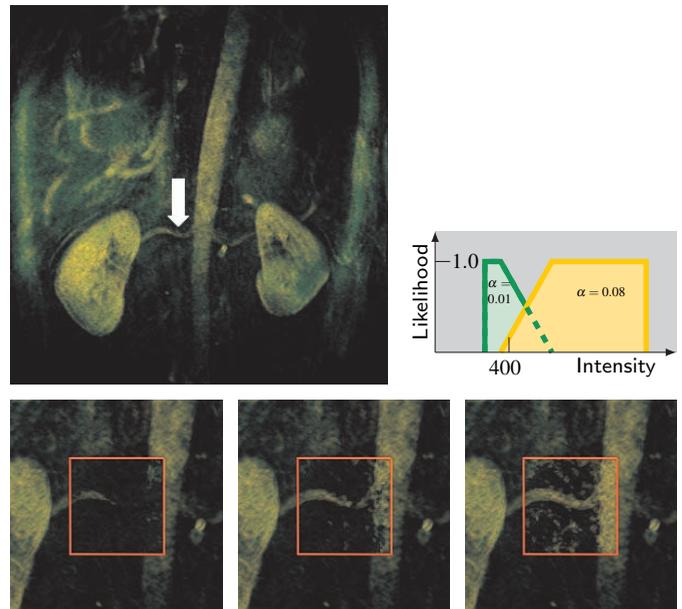


Fig. 10. The sensitivity lens applied to an MR renal angiography. Top left: A traditional rendering, a suspected stenosis is denoted by the arrow. Bottom row: The sensitivity lens applied to the suspicious region reveals that there is no stenosis. The square in the middle shows uncertainty animation in material sync mode while the other parts employ a traditional rendering. Top right: The probabilistic TF.

Table 2. FPS Performance of Uncertainty Animation

Data set Size	MR renal 512x512x56	CT neck 512x512x512	MR brain 432x432x160
Traditional Animation	42.7	9.2	31.0
	38.4	8.2	28.4

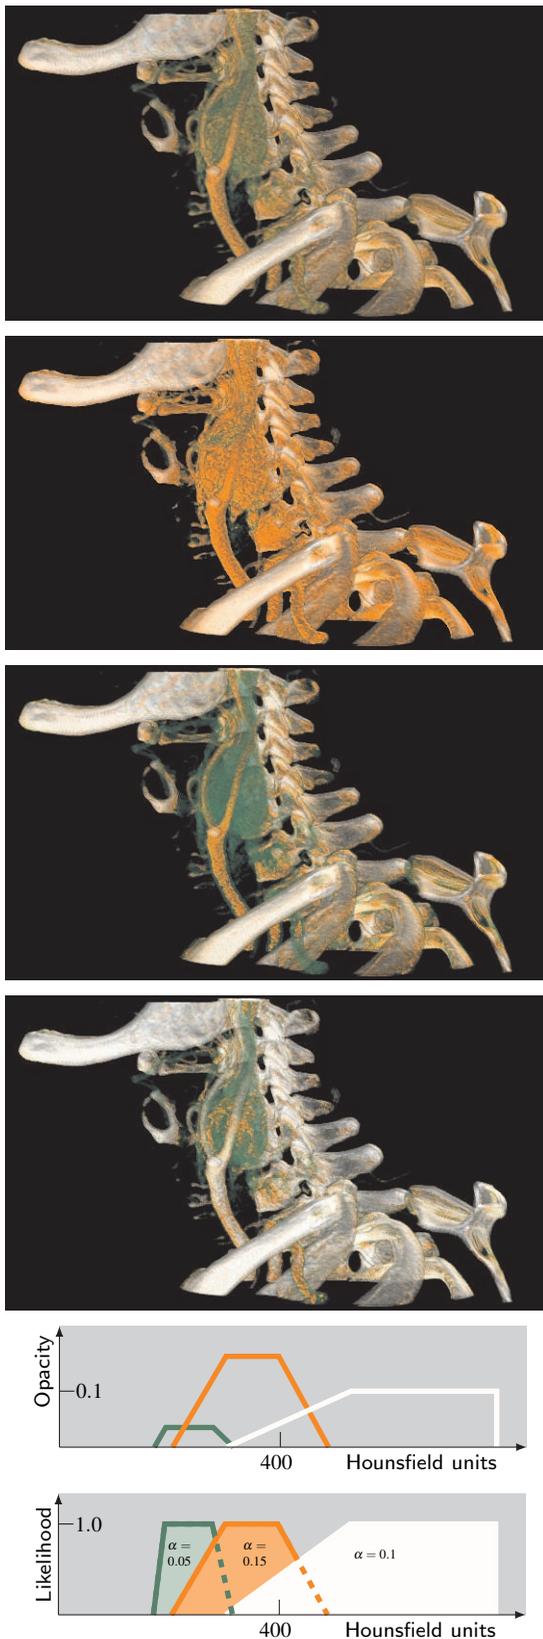


Fig. 11. CT examination of thyroid tumor. The top image is a traditional diagnostic visualization with overlapping TF components for tumor (green), vessels (orange), and bone (white). The following three images are different frames from an uncertainty animation in material sync mode. The dramatic differences show that there is much inherent uncertainty in the traditional rendering to be explored. The traditional and probabilistic TFs are given, the latter at the bottom.

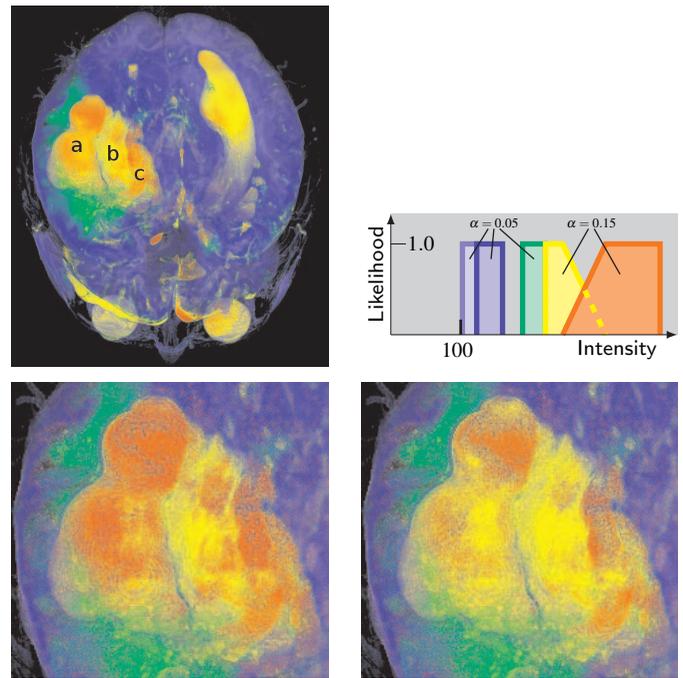


Fig. 12. MR brain examination. There is a tissue complex in the right brain hemisphere consisting of a cyst (a), a hemorrhage (b), and a ventricle (c), all three in the yellow-orange span. The complex is surrounded by an edema (green). The diagnostic task is to determine the borders between the cyst and the hemorrhage. Top left: Traditional rendering. Top right: The probabilistic TF. Bottom row: Two zoom-ins of uncertainty animation frames in random mode.

5 Conclusions

In this paper probabilistic animation methods have been presented as an effective approach for uncertainty visualization. The foundation is an explicitly probabilistic formulation of the TF, having benefits also for non-animation rendering. The proposed uncertainty animation can effectively expand the decision support for the physician running the visualization, as opposed to static, manually tuned renderings where important alternative presentations may be left out.

A disadvantage of animation is that the movement and flickering in the image causes visual fatigue. The proposed sensitivity lens application should, however, not suffer from this as the animation is limited to short periods and small regions. The fact that the animation speed cannot be greater than the overall frames per second in the rendering may be a limiting factor in some cases. An interesting area for future work is the representation of probability in the time domain, since perceived probability may be non-linearly mapped to the occupation ratio in the animation cycle.

The experimental study shows that for a fixed TF setting, the uncertainty animation in material sync mode is clearly more efficient for stenosis assessment than a static rendering. In fact, it comes quite close to the “gold standard” of free manual adjustment, a clear indication that time-consuming TF interaction can to some extent be replaced by uncertainty animation. Even though the intentionally poor conditions of the simulation would not be acceptable in a clinical situation, the physicians considered the test to have bearing on real diagnostic work.

Even though the presented methods stem from studies of clinical diagnostic work, further studies are needed to assess the full impact of the proposed methods in a real clinical situation. An interesting question is the generality, whether the clinical usefulness of the uncertainty animation is limited to the sensitivity lens application or to angiographies. The informative renderings in figures 11 and 12 indicate that the proposed method indeed can provide benefits in a wide range of cases. Another topic for future work is to develop an intuitive user control of the animation parameters.

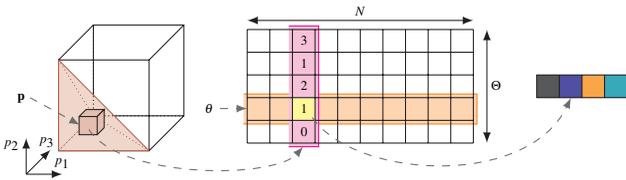


Fig. 13. Probabilistic animation in the general case: a volume classified with respect to M materials, exemplified for $M = 3$. From left to right: The material probability vector $\mathbf{p} = (p_1, \dots, p_M)$ is looked up in the material probability LUT, restricted by $\sum_{m=1}^M p_m \leq 1.0$. The found value is a column index into the $\Theta \times N$ material selection LUT. Together with the animation index θ , a material index is looked up and finally used to find the resulting sample color.

Acknowledgements

The authors thank Monica Tavanti, Nils Dahlström, and Örjan Smedby for assistance in the experimental study design and evaluation. Thanks also to Johan Kihlberg and Petter Quick at CMIV for data set retrieval and to the physicians in the user study. This work has been funded by the Swedish Research Council, grant 621-2003-6582, and the Swedish Foundation for Strategic Research through the Strategic Research Center MOVIII. There are pending patents based on this work.

A Generalization of Probabilistic Animation

The implemented probabilistic animation scheme is based on a classification in form of a 1D intensity-based TF. The relevant starting point for an arbitrary classification scheme is a volume with material probabilities. An outline of the probabilistic animation for the general case is presented in figure 13. A sample from the classified volume corresponds to a location in an M -dimensional material probability LUT. The material probability LUT points into a material selection LUT and the sample color is then found in the material color array.

The sizes of the LUTs are not expected to be a problem. Judging from the experience with the TF-based scheme, a size of 10^M would be sufficient for the material probability LUT and can probably be reduced further for large M . Moreover, the restriction that the sum of probabilities is less than 1.0 allows for reduced storage. The number of columns in the material selection LUT, N , equals the number of unique ways to fill Θ slots from a set of $M + 1$ material indices. This is known as the *multichoose* operation:

$$N(M, \Theta) = \binom{(M+1) + \Theta - 1}{\Theta} = \frac{(M+\Theta)!}{M!\Theta!} \quad (9)$$

$$N(4, 10) = 286 \quad N(8, 10) = 43758 \quad N(4, 20) = 63756$$

Thus, this LUT size is not a problem for reasonable values of M and Θ .

B Details of Experimental Study

The vessel model, shown in figure 8, was validated as realistic by three radiologists before the start of the study. In order to make the assessment challenging, the three materials (vessel lumen, vessel wall/plaque and background) were given highly overlapping intensity distributions. A realistic appearance of the vessel lumen was achieved by making the intensity decrease for increasing radius, causing the transition to the lower intensities of the wall and plaque to become very subtle. 24 different vessels were tested with randomized variations in stenosis degree, but also other parameters (radius and horizontal position of vessel, vertically varying horizontal offset of vessel, and vertical position of stenosis) were given varying values to make each test independent. Each vessel was given a randomized TF, consisting of a single trapezoid, designed to be more or less suboptimal. The transformation of eq. 5 was used to create comparable TFs for the methods. The TFs and the viewpoint were not adjustable.

Each physician assessed the stenosis of the vessels on a four-degree scale: 0-25%, 25-50%, 50-75%, and 75-100%. For the case that the stenosis was not assessable at all, the radiologist had a fifth possible outcome, "Unknown". The scale was indexed 0, 1, 2, 3 and the absolute index difference to the target index was used to measure accuracy. The "Unknown" response was given a fix error of 3.

In total, each test consisted of 72 trials, i.e. all method/vessel combinations. The order of the trials was determined by a controlled randomization. To avoid carry-over effects, the order in which the methods occurred was counterbalanced both within each subject's trial set and between the subjects. Moreover,

Table 3. Test Subject Characteristics

	1	2	3	4	5	6	7	8	9	10	11	12
Age	37	46	53	37	55	60	37	53	51	51	56	37
Sex	M	F	M	M	M	M	F	M	M	M	M	M
Years of experience	8	10	21	8	23	34	2	33	20	23	30	3

the controlled randomization ensured that two trials for a certain vessel was separated by at least 20 other trials. To familiarize the subjects with the task before the test, the setup was demonstrated by 12 dry-run trials (all method/stenosis degree combinations) where the TF could be adjusted and the accuracy of the assessment was fed back to the subject. The "gold standard" test with free manual TF adjustments was carried out for the 24 vessels after the main test.

The characteristics of the test subjects are given in table 3.

References

- [1] S. Djurcilov, K. Kim, P. F. J. Lermusiaux, and A. Pang. Volume rendering data with uncertainty information. In D. Ebert, J. Favre, and R. Peikert, editors, *Data Visualization '01*, pages 243–252, 355–356. Springer, 2001.
- [2] R. A. Drebin, L. Carpenter, and P. Hanrahan. Volume rendering. In *Proceedings of Computer Graphics and Interactive Techniques*, volume 22, pages 65–74. ACM SIGGRAPH, 1988.
- [3] C. R. Ehlschlaeger, A. M. Shortridge, and M. F. Goodchild. Visualizing spatial data uncertainty using animation. *Computers & Geosciences*, 23:387–395, May 1997.
- [4] K. Engel, M. Hadwiger, J. Kniss, C. Rezk-Salama, and D. Weiskopf. *Real-Time Volume Graphics*. A.K. Peters, Ltd, 2006.
- [5] N. D. Gershon. Visualization of fuzzy data using generalized animation. In *Proceedings of IEEE Visualization '92*, pages 268–273, 1992.
- [6] G. Grigoryan and P. Rheingans. Point-based probabilistic surfaces to show surface uncertainty. *IEEE Transactions on Visualization and Computer Graphics*, 10:564–573, 2004.
- [7] C. D. Hansen and C. R. Johnson, editors. *The Visualization Handbook*. Elsevier Butterworth-Heinemann, 2005.
- [8] T. Hengl. Visualisation of uncertainty using the HSI colour model: computations with colours. In *Proceedings of the 7th International Conference on GeoComputation*, pages 8–17, 2003.
- [9] C. Johnson. Top scientific visualization research problems. *IEEE Computer Graphics and Applications*, 24(4):13–17, July/August 2004.
- [10] G. Kindlmann, R. Whitaker, T. Tasdizen, and T. Möller. Curvature-based transfer functions for direct volume rendering: Methods and applications. In *Proceedings IEEE Visualization 2003*, pages 513–520, 2003.
- [11] J. M. Kniss, R. V. Uiter, A. Stevens, G.-S. Li, T. Tasdizen, and C. Hansen. Statistically quantitative volume visualization. In *Proceedings IEEE Visualization 2005*, pages 287–294, 2005.
- [12] C. Lundström, P. Ljung, and A. Ynnerman. Extending and simplifying Transfer Function design in medical Volume Rendering using local histograms. In *Proceedings EuroGraphics/IEEE Symposium on Visualization 2005*, pages 263–270, 2005.
- [13] C. Lundström, P. Ljung, and A. Ynnerman. Local histograms for design of Transfer Functions in Direct Volume Rendering. *IEEE Transactions on Visualization and Computer Graphics*, 12:1570–1579, Nov/Dec 2006.
- [14] A. T. Pang, C. M. Wittenbrink, and S. K. Lodha. Approaches to uncertainty visualization. *The Visual Computer*, 13:370–390, November 1997.
- [15] A. Persson, N. Dahlström, L. Engellau, E.-M. Larsson, T. Brismar, and Ö. Smedby. Volume rendering compared with maximum intensity projection for magnetic resonance angiography measurements of the abdominal aorta. *Acta Radiologica*, 45:453–459, 2004.
- [16] P. Rheingans and S. Joshi. Visualization of molecules with positional uncertainty. In E. Gröller, H. Löffelmann, and W. Ribarsky, editors, *Data Visualization '99*, pages 299–306. Springer-Verlag Wien, 1999.
- [17] P. J. Rhodes, R. S. Laramée, R. D. Bergeron, and T. M. Sparr. Uncertainty visualization methods in isosurface rendering. In M. Chover, H. Hagen, and D. Tost, editors, *Eurographics 2003, Short Papers*, pages 83–88, 2003.
- [18] T. Zuk and S. Carpendale. Theoretical analysis of uncertainty visualizations. In R. F. Erbacher, J. C. Roberts, M. T. Gröhn, and K. Börner, editors, *Visualization and Data Analysis 2006*, volume 6060.