Eurographics Workshop on Visual Computing for Biology and Medicine (2014)
I. Viola, K. Bühler, and T. Ropinski (Editors)

Evaluation of Transfer Function Methods in Direct Volume Rendering of the Blood Vessel Lumen

G. Läthén[1,3], S. Lindholm[1,3], R. Lenz[1,3], and M. Borga[2,3]

[1]Department of Science and Technology, Linköping University
[2]Department of Biomedical Engineering, Linköping University
[3]Center for Medical Image Science and Visualization, Linköping University Hospital

Abstract

Visualization of contrast enhanced blood vessels in CT angiography data presents a challenge due to varying concentration of the contrast agent. The purpose of this work is to evaluate the correctness (effectiveness) in visualizing the vessel lumen using two different 3D visualization strategies, thereby assessing the feasibility of using such visualizations for diagnostic decisions. We compare a standard visualization approach with a recent method which locally adapts to the contrast agent concentration. Both methods are evaluated in a parallel setting where the participant is instructed to produce a complete visualization of the vessel lumen, including both large and small vessels, in cases of calcified vessels in the legs. The resulting visualizations are thereafter compared in a slice viewer to assess the correctness of the visualized lumen. The results indicate that the participants generally overestimated the size of the vessel lumen using the standard visualization, whereas the locally adaptive method better conveyed the true anatomy. The participants did find the interpretation of the locally adaptive method to be less intuitive, but also noted that this did not introduce any prohibitive complexity in the work flow. The observed trends indicate that the visualized lumen strongly depends on the width and placement of the applied transfer function and that this dependency is inherently local rather than global. We conclude that methods that permit local adjustments, such as the method investigated in this study, can be beneficial to certain types of visualizations of large vascular trees.

Categories and Subject Descriptors (according to ACM CCS): Computer Graphics [I.3.6]: Methodology and Techniques—Computer Graphics [I.3.7]: Three-Dimensional Graphics and Realism—Computer Graphics [I.3.8]: Applications—

1. Introduction

Volumetric visualization techniques are now commonly available in the modern clinical workflow. There are two main arguments for their use. First, the increasing size of the data generated by modern imaging modalities makes traditional slice viewing time consuming. Second, volumetric visualizations of inherently three-dimensional structures often provide more intuitive overviews, and are also easier to interpret for medical personnel that lack extensive training in two-dimensional imaging. As such, three-dimensional techniques have found success in a number of collaborative scenarios such as preoperative planning of liver surgery [SPSP02], neurosurgical interventions [BHWB07] and surgical treatment of lung cancer [HM07] as well as full body virtual autopsies [DJV*06, LWP*06]. However, a major challenge is still the interaction with these techniques in order to produce a reliable and correct visualization of the anatomy. Therefore, the techniques still have limited use in applications which include diagnostic decisions based on the geometry of the anatomy, such as grading of vessel stenosis.

The purpose of this work is to evaluate and compare the correctness of two volume rendering setups for visualizing the blood vessel lumen. We compare a traditional method, which operates on a global scale, with a method which locally adjusts the visualization depending on contrast agent concentration (i.e. the lumen intensity). In the context of this work, the method referred to as "standard" uses a global intensity window which the user can translate and scale. The method referred to as "locally adjusted" uses an intensity window that can be adjusted by global translation (per-

DOI: 10.2312/vcbm.20141197

formed by the user) and local shifts (performed automatically) as described by Läthén et al. [LLL*12]. Both methods are described in Section 2.2.

2. Background

The methods in this study all use the traditional direct volume rendering technique (DVR) [DCH88]. This technique is based on tracing rays through the volume data, while at each point accumulating color and opacity attributes given by a mapping of the data values. This mapping, referred to as the transfer function (TF), is the main tool for controlling the visualization by assigning visual attributes to data values. The simplest type of TF is parameterized only on data values and is very similar to the windowing tool traditionally used in viewing medical images. That is, the typical use of the TF is to select a relevant data value range, typically aimed at a specific tissue type, and to maximize the visibility of that range. For computed tomography (CT) data, the values are calibrated by the Hounsfield scale which makes it relatively easy to design TFs targeting particular tissue types. To save time, many DVR systems include preset TFs which can be directly applied to different studies, e.g. bone fractures, soft tissue studies etc. Blood vessels however, require the injection of a contrast agent to achieve a useful level of contrast to the surrounding soft tissue. The observed concentration of the contrast agent in the blood stream is dependent on a number of factors, such as the timing between injection and scan, the blood flow and potential vessel pathologies. The result is that the observed value range representing the vessels is highly uncertain. When applying a preset TF for CT angiography (CTA), the user needs to carefully assess the visualization with two objectives; that the TF correctly aligns with the true value range of the lumen content (accuracy), and that parts of non-relevant tissues do not obstruct the view of the vascular tree (usefulness).

2.1. Related Work

The visualization process can be made more robust by adding additional parameters to the TF, such as derivatives of the data or structural measures [HKG00, KD98, KWTM03, Lev88, SWB*00]. However, adding more parameters makes the TF design challenging. To overcome this, different high-level abstractions and specialized user widgets have been suggested [RSKK06, KKH02, TPD06, VHHFG05, GMY11]. These ideas have not yet reached the tools used in clinical practice, and for this reason we restrict the TF models in the study to be defined over a single dimension (the gray scale data value) which can be intuitively interpreted by the end user.

Regarding vessel visualization in particular, a majority of current methods use a prior segmentation as input to the visualization. See overviews [BFC03, PO08, KGNP12]. Most common is the use of vessel centerlines to guide curved planar reconstruction and its variations [GOH*10, KFW*02, AMB*13, MMV*13, MVB*12, KPS14]. In daily clinical routine, we consider such segmentations too time consuming and does not use this additional prior information in the methods evaluated in this study.

2.2. Transfer function methods

In CTA, there are many situations where the TF preset is misaligned relative to the value range of the lumen content and thus need to be adapted. For example, when the contrast agent concentration was lower than expected (causing the value range to shift globally) or if the contrast agent did not reach all parts of the vascular tree due to a stenosis (causing the value range to shift locally). Therefore, a TF method needs certain degrees of freedom (DoF) so that the user can interact with, and adapt, the visualization to the data. These DoFs should not be confused with additional TF parameters presented in the previous section - they rather introduce freedom in changing the shape and position of the initial TF preset.

This study compares two transfer function methods aimed at visualizing vessels. Both methods contain a single visible graphical widget used to adjust the transfer function. Additionally both methods also include a ramp widget (hidden to the user) which highlights calcified plaques and bone as white structures for increased context information. This widget remained fixed across all evaluations.

The first method is referred to as "standard" and closely resembles what is commonly available in clinical workstations. The TF is represented as a trapezoid shaped widget as exemplified in Figure 1 (bottom, left), along with the output rendering in Figure 1 (top, left). To interact with the visualization, the user can translate and change the width of the TF using a mouse interface.

The second method is referred to as "locally adjusted" and is based on the same trapezoidal TF widget but includes an additional spatial dependency that, when evaluated, automatically shifts the widget (in the value domain) depending on characteristics of the data where it is applied (in the spatial domain). The spatial dependency is pre-computed, and designed with the goal to reveal as much vessel structures as possible at all locations by adjusting to the (spatially) varying intensity levels of the lumen content. This is done by finding the TF shift which results in maximum vesselness according to the definition by Sato et al. [SWB*00]. The computation is based on a number of application dependent parameters (scale and intensity value ranges) which in general can be fixed for a particular anatomy. Thus, in a clinical setting, the method can be fully automatized in a pre-processing step. Note that the original work [LLL*12] does not investigate any scheme for interacting with the locally shifted TFs during visualization. In this work, we propose to illustrate the method's TF as shown in Figure 1 (bottom, right), along with the output rendering in Figure 1 (top,

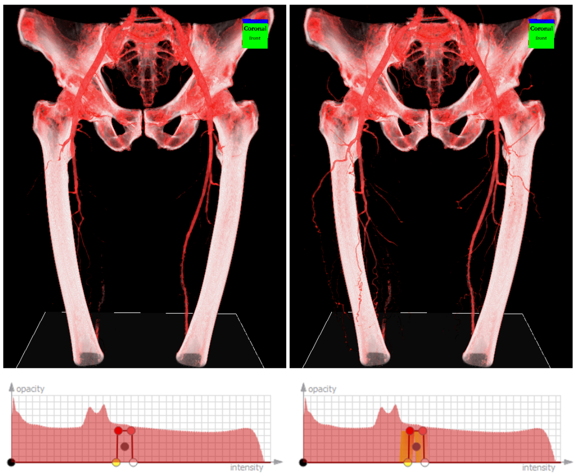

Figure 1: *The views in the test application. The left part shows the output of the rendering (top) using the "standard" method with a trapezoidal transfer function (bottom), while the right part displays the output and transfer functions of the "locally adjusted" method. The user can interact with both transfer functions and navigate in the data by zooming, panning and rotating.*

right). The maximal spatially dependent shift is represented in the interface as a yellow shadow overlay (i.e. the left-most part of the yellow shadow represents TF widgets which are maximally shifted towards lower values).

An example where the TF is misaligned with regards to the value range of the lumen is shown in Figure 1 where the "standard" method (left) fails to visualize the small vessels. The "locally adjusted" method (right) has shifted the TF to better visualize these small vessels. To further interact with the visualizations and verify that the TF is well adapted to the vessels' value range, both methods possess two DoFs. Global shifts are handled identically for both methods by translations of the TF widgets (1st DoF both methods), while local shifts are handled differently between the two methods. The "standard" TF method does not permit local adjustments, so local value range shifts can only be addressed by scaling the widgets to accommodate the globally widened value range (2nd DoF "standard" method). The "locally adjusted" method, on the other hand, allows local adjustments to be computed and applied automatically. The maximum strength of this automatic local shift can be regulated on a global level as a linear range from no local adjustments to full local adjustments (2nd DoF "locally adjusted" method). In the interface, the second DoF is directly related to the size of the yellow shadow of the TF widget.

3. Materials and Methods

The study was performed as a within subject evaluation (all subjects are exposed to all test cases), consisting of an introduction followed by three separate sequential parts in; 1) task completion, 2) result verification and 3) a questionnaire with ranked questions. At the conclusion of the questionnaire phase, additional follow-up questions were asked regarding notable answers or outliers.

This section will first detail the aim of the study and the evaluated variables before providing details to each part of of the trial.

3.1. Requirements specification

To define what should be assessed in the study, a set of requirements for an effective visualization method were defined:

R1. The coupling between the interaction and its effect on the visualization should be intuitive
R2. The interaction should support the user in the task to explore the vessel structures in the visualization
R3. It should be easy to make a clinical interpretation of the visualization (different anatomy and their structure should be clear)
R4. The visualization should only contain relevant anatomy
R5. The method should result in a safe (correct) diagnosis

We note here that a visualization method also must be efficient in order to be successful (indicating a low time consumption), but repeat that we in this study focus only on effectiveness. Time consumption was thus not explicitly targeted in the questionnaire and no timings were performed during task completion. Furthermore, the study was designed such that the two methods were evaluated in parallel during the task completion to ensure that the participants felt that they had achieved as-good-as-possible results with both methods. This design, together with the participants' previous familiarity with one of the methods, further prohibited reliable timings.

3.2. Data and participants

The type of medical cases used in the study are three-dimensional CTA of calcified vessels in the leg, acquired using a Siemens SOMATOM Definition Flash CT Scanner (Siemens, Erlangen, Germany). Prior to the study, three unique cases were anonymized and pre-processed to remove the CT patient table and cropped to provide a simple and interactive test application setup. Five radiologists (one woman, four men, ages 33-63 years) with different levels of experience were recruited for the study. Their clinical radiological experience varied from 2 to 30 years, and their 3D visualization experience ranged from "minor" to "extensive" (one participant had partly technical background). Please note that, due to the highly specialized profile required of the participants, this corresponds to approximately 50% of the eligible work force at the hospital where the study was performed. Since one case was used for introduction to the methods, the five radiologists ranked two cases, resulting in ten replies for the analysis.

3.3. Test application setup and interface

For each case, the participant was presented with two views of the current dataset, with renderings using the different TF methods presented in Section 2.2. The left part of Figure 1 shows the visualization using the "standard" TF method, whereas the right part shows the "locally adjusted" TF. The interface allows panning, zooming and rotating both views in sync, in addition to interacting with the TF as described in Section 2.2.

Note that the evaluations of the two methods for a single case were performed in parallel (as opposed to sequentially). This design was chosen to minimize the learning effects that occur within each case (familiarity with the data) and to allow the participant to alternate between the methods in order to find the best possible image for each method.

After the TF interactions were completed, the correctness of the visualizations was assessed using slice views of the data and the output of the TFs, as in Figure 2.

3.4. Questionnaire

To measure the degree of fulfilling the requirements in Section 3.1, a questionnaire was designed with the following questions, ranked from 1 to 5. The phrasing in this paper is a translation from the original questionnaire in the native language of the radiologists. The interpretation of the ranks are specified after each question. In addition, the requirement targeted by each question is noted in bold face:

1. *Interaction feedback*: How did you experience the direct feedback between the visualization and your interaction? (1=unclear, 5=clear) (**R1**)
2. *Goal achievement*: To what extent could you attain the image you wanted? (1=low extent, 5=high extent) (**R2**)
3. *Interaction simplicity*: How did you experience the process of attaining the image you wanted? (1=hard, 5=easy) (**R2**)
4. *Anatomical identification*: How did you experience the visualization in terms of identifying different anatomies and their structure? (1=unclear, 5=clear) (**R3**)
5. *Anatomical relevance*: To what extent did the visualization contain irrelevant anatomy? (1=high extent, 5=low extent) (**R4**)

In addition, the questionnaire included two questions with binary response:

6. *Anatomical interference*: Did the irrelevant anatomy interfere with the task of visualizing the blood vessels? (yes, no) (**R4**)
7. *Visualization correctness*: When comparing with axial slice images, which method gives the most correct visualization of the vessels' lumen? ("standard", "locally adjusted") (**R5**)

3.5. Procedure

The procedure was constructed as an information and training phase (using one of the three datasets) followed by the task completion and evaluation of two cases (using the two remaining two datasets in random order).

Initially, the participant was informed that the objective would be to create visualizations of the vessel lumen over the complete vascular tree by interacting with the transfer functions. Then the participant read a printed text describing the background, the tasks and a description of the user interface. The application was then demonstrated using one case such that the participant could familiarize him/her-self with the tools. Questions regarding the interface, including the function and interpretation of its parts, were clarified at this point until the participant was confident in understanding the task and the interface.

During the task completion phase, one of the two remaining datasets was loaded and the participant was assigned to start with one of the two TF methods (randomly selected). When the result was satisfactory, the participant switched method to solve the same task with the second TF method. This step was performed using the side-by-side interface illustrated in Figure 1. Before finalizing the task, the participant was allowed to iterate between methods and views at will to tune both results. Finally, when the participant was satisfied, all further TF interaction was locked and the instructor opened the slice viewer, illustrated in Figure 2.

After the task completion phase, the participant was assigned to assess which of the methods gave the most correct visualization of the lumen in comparison with the slice data. This was performed by inspecting the slice views in Figure 2, which show the slice data and output overlays of the two TF methods. The participant examined a stack of slices in assigned regions where the output of the methods conflict. The participant was also allowed to inspect regions of further interest, such as suspected stenoses etc. When this step was completed, the questionnaire was answered, and the procedure was repeated with the second, and final, task dataset. Finally, a short informal discussion took place where the participant was encouraged to more freely discuss the questions in the questionnaire or bring up other issues the user was experiencing. In the case of outlying answers, these were typically followed up at this stage with additional questions from the instructor. (Through the remainder of the paper, these follow-up discussions will be collectively referred to as "the discussions").

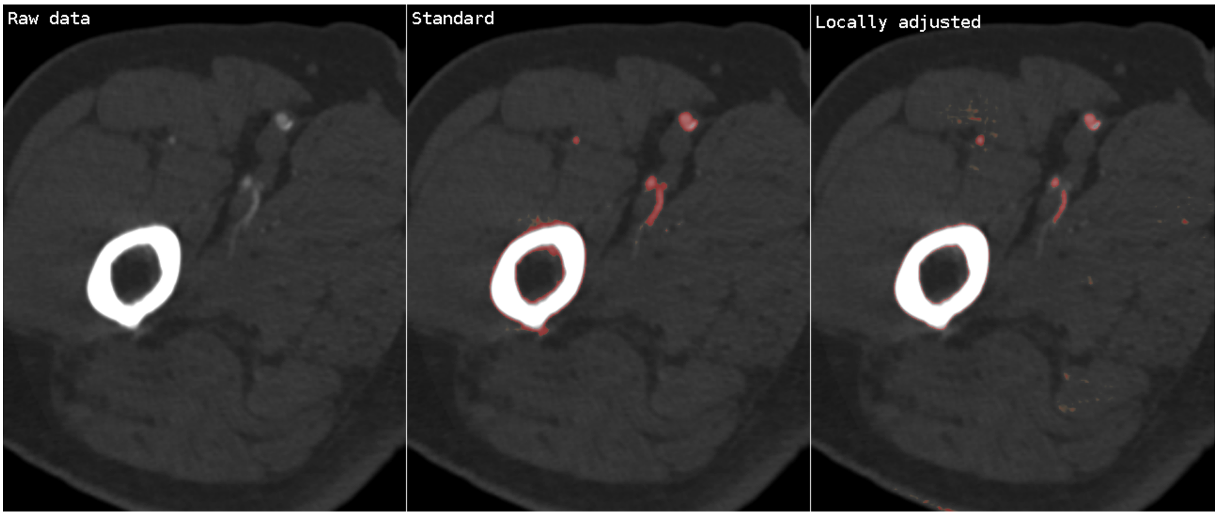

Figure 2: *Slice views of* (left) *the data,* (middle) "standard" method and *(right)* "locally adjusted" method. The "standard" and "locally adjusted" method views depict the data overlaid with the classification result from the respective method. The three views were used by the radiologists to assess to which extent the two methods produced clinically correct results.

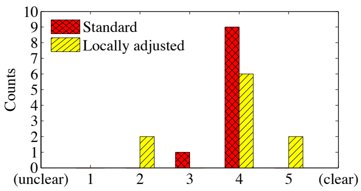

(a) 1. Interaction feedback

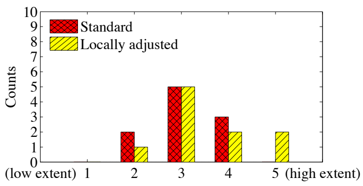

(b) 2. Goal achievement

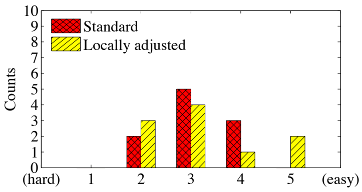

(c) 3. Interaction simplicity

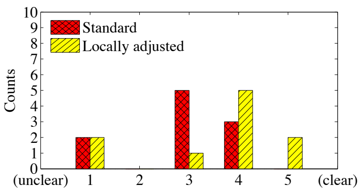

(d) 4. Anatomical identification

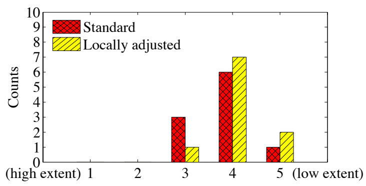

(e) 5. Anatomical relevance

Figure 3: *Frequency plots of the ranked responses to questions 1-5 in the questionnaire (Section 3.4) for the two methods "standard" and "locally adjusted".*

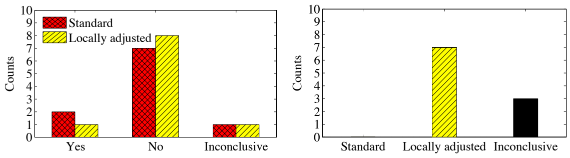

(a) 6. Anatomical interference (b) 7. Visualization correctness

Figure 4: *(a): The number of "yes/no" replies to question 6 in the questionnaire (Section 3.4) for the respective methods. (b): The number of votes for the respective methods "standard" and "locally adjusted" in response to question 7 in the questionnaire (Section 3.4).*

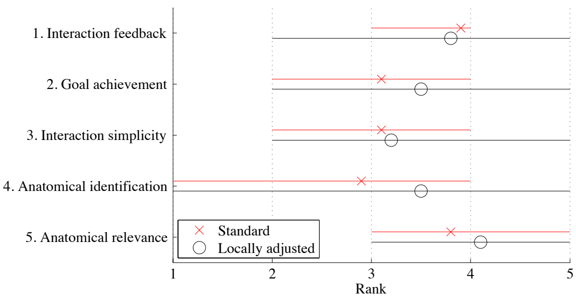

Figure 5: *Average values for the ranked responses to questions 1-5 in the questionnaire (Section 3.4) for the two methods "standard" and "locally adjusted" (higher rank generally corresponds to "better"). The markers show the average values, while the horizontal lines show the minimum and maximum ranks for each question.*

4. Results

The results from 10 individually evaluated cases were collected (two cases per participant), see Table 1. Answer frequency plots for the ranked questions 1-5 are shown in Figure 3 for the respective methods. For the binary response questions 6 and 7, frequency plots are shown in Figure 4. Any inconclusive number of replies within single questions are due to missing or indecisive replies and are further discussed in Section 5. The average, and min/max, scores for answers to ranked questions 1-5 are shown in Figure 5 for the two methods. Note that computing averages on an ordinal scale can be discussed, but we include this figure as an illustration to better convey the *trends* of the full set of results.

Note also that, due to the low number of participants, we have chosen not to perform a full statistical evaluation. We provide instead all the collected data in full, and discuss trends and outliers together with the comments and observations from individual participants collected during the post-questionnaire discussions.

5. Discussion

We will start by discussing each question in turn. For simplicity, the post-questionnaire discussions and follow-up questions will, in this section, be referred to as "the discussions":

1. Interaction feedback: This question targets requirement R1 (Section 3.1) which states that the interaction and its effect on the visualization should be intuitive, so that the user feels in control over the visualization. Since the "standard" method is based on a simpler interaction model, the expectation was that this method would score superior to the "locally adjusted" method in this category. This was partly indicated by the results with a slightly higher average rating for the "standard" method (Figure 5). The detailed frequency plot (Figure 3(a)) also indicates that the participants agreed more on the ranking of the "standard" method (lower variance) whereas the level of understanding of the "locally adjusted" method differed more between participants. This was also evident during the introduction of the methods, as more time was generally required for the instructor to explain the idea behind the "locally adjusted" method and the interpretation of its free parameters. We believe this to be an expected result since most participants have previous experience with a variation of the "standard" method through the current on-site software. The average score for the "locally adjusted" method was still at an acceptable level, and the participants did not express concern with the concept after practicing.

2. Goal achievement: This question targets requirement R2 (Section 3.1) which states that the method should support the user when exploring the vessel structures. In the test setup, this question measures how well the method performs with the specific goal of visualizing the lumen for all parts of the vascular tree. Since the "locally adjusted" method was designed to support this task, the expectation was that this method would be superior in this case. This was also supported by the participants by a higher average ranking (Figure 5), albeit with a slightly wider distribution of answers (Figure 3(b)). The higher ranking was also supported during following discussions with the participants where they generally expressed the feeling that the "locally adjusted" method better succeeded in representing large and small vessels simultaneously.

3. Interaction simplicity: Like the previous question, this question also targets requirement R2, but from the perspective of task complexity. This question measures how easy/hard it is to reach the final image with the underlying goal to visualize all parts of the vascular tree. The rankings in Figure 5 and Figure 3(c) state that the methods performed quite similar in this case with a slightly better ranking for the "locally adjusted" method. This verifies that the conceptually more complex "locally adjusted" method is still within an acceptable complexity level for the participants in this study.

4. Anatomical identification: This question targets re-

Question	1		2		3		4		5		6		7	
Participant	Std	LA	Std	LA	Std	LA	Std	LA	Std	LA	Std	LA	Std	LA
1	4	5	3	5	3	5	3	5	4	5	No	No		x
	4	5	4	5	4	5	4	5	5	5	No	No		x
2	4	4	3	3	2	2	1	1	4	4	No	No	x	x
	4	4	3	4	4	4	1	1	4	4	No	No		x
3	3	4	3	3	3	3	3	4	4	4				x
	4	4	2	3	2	2	3	3	3	3	Yes	Yes		x
4	4	2	4	3	4	3	3	4	3	4	Minor	No		x
	4	2	3	3	3	3	3	4	3	4	No	No		x
5	4	4	2	2	3	2	4	4	4	4	No	No	x	x
	4	4	4	4	3	3	4	4	4	4	No	No	x	x

Table 1: *Table with complete set of results. The methods "standard" (Std) and "locally adjusted" (LA) were ranked by 5 participants on 2 medical cases.*

quirement R3 (Section 3.1) which states that the anatomy in the visualization can be clearly identified. In this case, the "locally adjusted" method ranks better than "standard" (Figure 5). In the discussions, it was noted that most participants considered the overall anatomical information, such as muscle tissue, bones etc., very similar between the two methods. The difference was mainly for the vascular structures which could appear quite different. For these structures, the "locally adjusted" method was considered as giving the clearest representation of the structure. The main difficulty when using the "standard" method was the visual separation of calcified plaques and the lumen, which often became merged. One participant (No. 2 in Table 1, 63 years, minor visualization experience) found it very hard to make any clinical judgment of the structure of the anatomy and ranked both methods as 1 (see Figure 3(d)). However, since the participant had minor prior experience with 3D visualizations, discussions revealed that more practice and experience could alleviate this issue.

5. Anatomical relevance: This question targets requirement R4 (Section 3.1) which states that the visualization should only contain relevant anatomy. The motivation behind this requirement is that too many irrelevant structures tend to occlude the important parts in 3D visualizations. For this question, the "locally adjusted" method achieved a higher ranking (Figure 5) with no clear outliers (Figure 3(e)). The main difference between the methods was that the "standard" method generally included more soft tissue (muscle) which appears when trying to visualize the small low-contrast vessels. For this task, most participants expressed the wish to have a bone removal function similar to one implemented in their current on-site workstation.

6. Anatomical interference: The purpose of this question was to follow-up the requirement R4 (i.e. whether any irrelevant anatomy did actually interfere with visualizing the vessels). However, since neither of the two methods produced visualizations that contained much non-significant anatomy in the first place (see previous question), the results from this question are less meaningful. The number of "no" replies (Figure 4(a)) still indicates that the soft tissue that did appear was not considered disturbing. (The "inconclusive" replies related to a case with a low degree of irrelevant anatomy so the participant left the question blank.)

7. Visualization correctness: This question targets requirement R5 (Section 3.1) which states that the visualization should result in a correct diagnosis (i.e. that the structure presented in the visualization should be correct). The participants concluded in most cases (7 out of 10, Figure 4(b)) that the "locally adjusted" method was giving the most correct representation of the vessel lumen. Generally, it was noted that the "standard" method overestimated the size of the lumen, including more of the transitional regions between the lumen and the surrounding soft tissue or plaques. An example of such a case is shown in Figure 6. However, three of the results were considered inconclusive, in the sense that both methods performed similarly. Two situations occurred: Either the "standard" method overestimated the lumen size, while the "locally adjusted" method underestimated the size. Or, one method was better in one part of the vascular tree, while the other was better in other parts. During the discussions, it was noted that it was very difficult to assess the correctness of the visualization from the 3D view, where both methods presented plausible vessel structures in most cases. One participant (7 years of experience in 3D visualization) made an initial guess for one case that the "standard" method presented the correct geometry by judging the 3D visualization. However, when validating the result in the slice views, it was concluded that "standard" had overestimated the lumen to a large extent.

Other: Several participants reacted positively on the goal of the assigned task (i.e. to create a visualization containing both large and small vessels). They confirmed that this is typically difficult due to the large variation of intensity values representing the vessel lumen. They also acknowledged that simultaneously visualizing both small and large vessels provides a valuable context. For example, a large number of

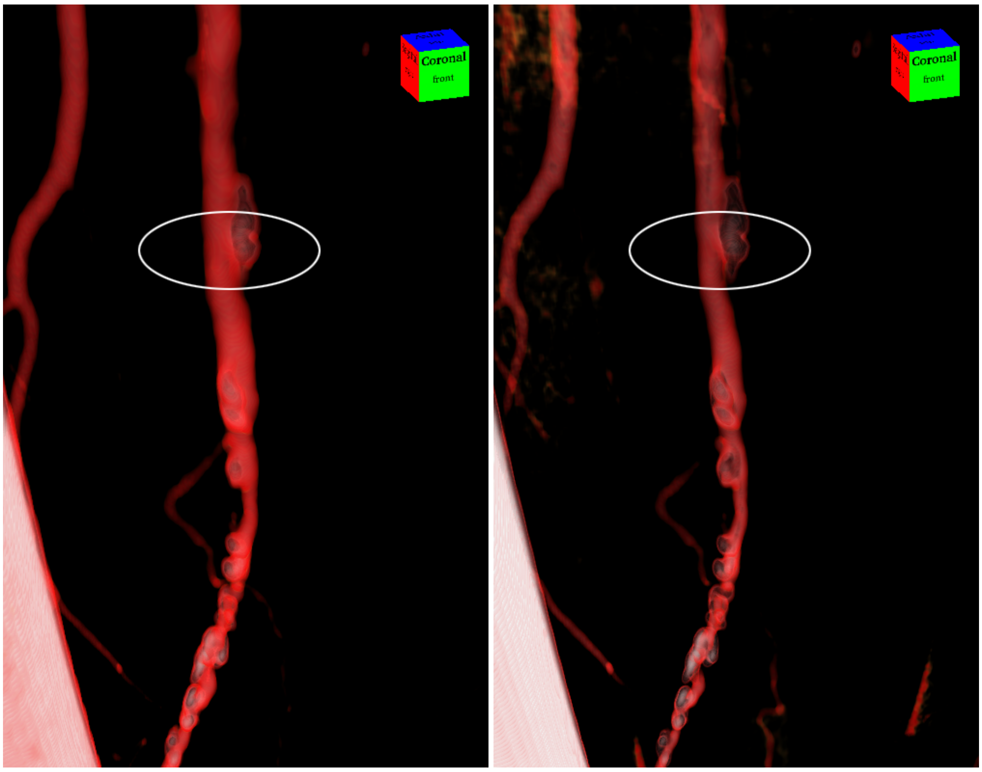

(a) Standard DVR (b) Locally adjusted DVR

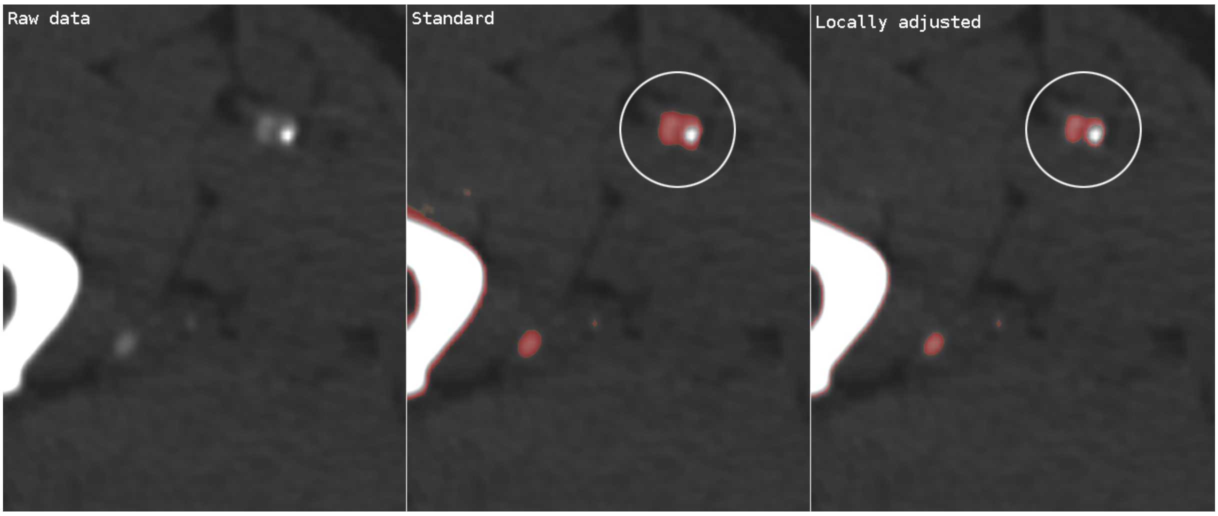

(c) Slice view

Figure 6: *Example where the "standard" method overestimated the lumen according to a participant. The circled regions in the figures correspond to the same structure. Figures (a) and (b) show the output from the visualization using the respective methods. Their output overlaid with the original data is shown in (c), middle and right.*

enlarged collateral arteries is an indication that a stenosis in the main artery has appeared gradually, rather than suddenly.

One participant (5) noted during the discussions that the "locally adjusted" method tended to converge faster to what was considered the best visualization. The feeling was that this method more quickly revealed the appropriate ranges in the data. However, since we did not focus on time consumption in this study (see Section 3.1) this claim remains to be studied in future work.

6. Conclusions and future work

With the goal to visualize the lumen of the entire vascular tree in the legs (i.e. both large and small vessels), the results indicate that the "standard" method generally overestimates the size of the lumen. This is connected to the fact that the width of the intensity window (TF) needs to be larger to include the intensity range of both small and large vessels. By this, we conclude that the width of the intensity window plays a crucial role in estimating the lumen size and needs to be carefully controlled for reliable vessel geometry assessment. Furthermore, the "locally adjusted" method did not introduce complexity for the user in terms of achieving a trustworthy image, even if it did have a more complex interpretation. Future work includes comparing measurements in 3D using the "locally adjusted" method with conventional angiography images as references. This also includes finding an optimal width of the TF preset to match the reference angiograms as close as possible. More work is also required in investigating the best way to visualize calcified plaques without occluding the vessel lumen.

Acknowledgments

We gratefully acknowledge Dr. Camilla Forsell (Linköping University) for valuable feedback on the user study design, and MD. Nils Dahlström and Prof. Anders Persson (Linköping University Hospital) for medical insights.

References

[AMB*13] Auzinger T., Mistelbauer G., Baclija I., Schernthaner R., Kochl A., Wimmer M., Gröller M., Bruckner S.: Vessel visualization using curved surface reformation. *IEEE Transactions on Visualization and Computer Graphics 19*, 12 (Dec. 2013), 2858–2867. doi:10.1109/TVCG.2013.215. 2

[BFC03] Bühler K., Felkel P., Cruz A. L.: Geometric methods for vessel visualization and quantification – a survey. In *Geometric Modelling for Scientific Visualization*, Mathematics and Visualization. Springer Berlin Heidelberg, 2003, pp. 399–420. doi:10.1007/978-3-662-07443-5_24. 2

[BHWB07] Beyer J., Hadwiger M., Wolfsberger S., Bühler K.: High-quality multimodal volume rendering for preoperative planning of neurosurgical interventions. *IEEE Transactions on Visualization and Computer Graphics 13*, 6 (2007), 1696–1703. doi:10.1109/TVCG.2007.70560. 1

[DCH88] Drebin R. A., Carpenter L., Hanrahan P.: Volume rendering. In *Computer Graphics and Interactive Techniques (SIGGRAPH '88)* (New York, USA, 1988), ACM, pp. 65–74. 2

[DJV*06] Dirnhofer R., Jackowski C., Vock P., Potter K., Thali M. J.: Virtopsy: Minimally invasive, imaging-guided virtual autopsy. *RadioGraphics 26*, 5 (2006), 1305–1333. doi:10.1148/rg.265065001. 1

[GMY11] Guo H., Mao N., Yuan X.: WYSIWYG (what you see is what you get) volume visualization. *IEEE Transactions on Visualization and Computer Graphics 17*, 12 (Dec 2011), 2106–2114. doi:10.1109/TVCG.2011.261. 2

[GOH*10] Glasser S., Oeltze S., Hennemuth A., Kubisch C., Mahnken A. H., Wilhelmsen S., Preim B.: Automatic transfer function specification for visual emphasis of coronary artery plaque. *Computer Graphics Forum 29*, 1 (2010), 191–201. doi:10.1111/j.1467-8659.2009.01590.x. 2

[HKG00] Hladuvka J., König A., Gröller E.: Curvature-based transfer functions for direct volume rendering. In *Spring Conference on Computer Graphics* (2000), vol. 16, pp. 58–65. 2

[HM07] Hu Y., Malthaner R. A.: The feasibility of three-dimensional displays of the thorax for preoperative planning in the surgical treatment of lung cancer. *European Journal of Cardio-Thoracic Surgery 31*, 3 (2007), 506–511. doi:10.1016/j.ejcts.2006.11.054. 1

[KD98] Kindlmann G., Durkin J. W.: Semi-automatic generation of transfer functions for direct volume rendering. In *IEEE Symposium on Volume Visualization* (New York, NY, USA, 1998), ACM, pp. 79–86. doi:10.1145/288126.288167. 2

[KFW*02] Kanitsar A., Fleischmann D., Wegenkittl R., Felkel P., Gröller M.: CPR - curved planar reformation. In *IEEE Visualization, 2002. VIS 2002* (Nov. 2002), pp. 37–44. doi:10.1109/VISUAL.2002.1183754. 2

[KGNP12] Kubisch C., Glasser S., Neugebauer M., Preim B.: Vessel visualization with volume rendering. In *Visualization in Medicine and Life Sciences II*, Linsen L., Hagen H., Hamann B., Hege H.-C., (Eds.), Mathematics and Visualization. Springer Berlin Heidelberg, 2012, pp. 109–134. doi:10.1007/978-3-642-21608-4_7. 2

[KKH02] Kniss J., Kindlmann G., Hansen C.: Multi-dimensional transfer functions for interactive volume rendering. *IEEE Transactions on Visualization and Computer Graphics 8*, 3 (2002), 270–285. doi:http://doi.ieeecomputersociety.org/10.1109/TVCG.2002.1021579. 2

[KPS14] Kretschmer J., Preim B., Stamminger M.: Bilateral depth filtering for enhanced vessel reformation. In *IEEE/Eurographics Symposium on Visualization (EuroVis - Short Papers)* (2014), pp. 13–17. doi:10.2312/eurovisshort.20141150. 2

[KWTM03] Kindlmann G., Whitaker R., Tasdizen T., Möller T.: Curvature-based transfer functions for direct volume rendering: Methods and applications. In *IEEE Visualization 2003. VIS 2003* (Washington, DC, USA, 2003), pp. 513–520. doi:10.1109/VISUAL.2003.1250414. 2

[Lev88] Levoy M.: Display of surfaces from volume data. *IEEE Computer Graphics and Applications 8*, 3 (may 1988), 29–37. 2

[LLL*12] Lathen G., Lindholm S., Lenz R., Persson A., Borga M.: Automatic tuning of spatially varying transfer functions for blood vessel visualization. *IEEE Transactions on Visualization and Computer Graphics 18*, 12 (Dec. 2012), 2345–2354. 2

[LWP*06] LJUNG P., WINSKOG C., PERSSON A., LUNDSTRÖM C., YNNERMAN A.: Full body virtual autopsies using a state-of-the-art volume rendering pipeline. *IEEE Transactions on Visualization and Computer Graphics 12*, 5 (2006), 869–876. 1

[MMV*13] MISTELBAUER G., MORAR A., VARCHOLA A., SCHERNTHANER R., BACLIJA I., KÃŰCHL A., KANITSAR A., BRUCKNER S., GRÃŰLLER E.: Vessel visualization using curvicircular feature aggregation. *Computer Graphics Forum 32*, 3pt2 (2013), 231–240. doi:10.1111/cgf.12110. 2

[MVB*12] MISTELBAUER G., VARCHOLA A., BOUZARI H., STARINSKY J., KOCHL A., SCHERNTHANER R., FLEISCHMANN D., GROLLER M., SRAMEK M.: Centerline reformations of complex vascular structures. In *IEEE Pacific Visualization Symposium (PacificVis)* (Feb. 2012), pp. 233–240. doi:10.1109/PacificVis.2012.6183596. 2

[PO08] PREIM B., OELTZE S.: 3d visualization of vasculature: an overview. In *Visualization in medicine and life sciences*, Mathematics and Visualization. Springer Berlin Heidelberg, 2008, pp. 39–59. doi:10.1007/978-3-540-72630-2_3. 2

[RSKK06] REZK-SALAMA C., KELLER M., KOHLMANN P.: High-level user interfaces for transfer function design with semantics. *IEEE Transactions on Visualization and Computer Graphics 12*, 5 (2006), 1021–1028. doi:10.1109/TVCG.2006.148. 2

[SPSP02] SELLE D., PREIM B., SCHENK A., PEITGEN H.-O.: Analysis of vasculature for liver surgical planning. *IEEE Transactions on Medical Imaging 21*, 11 (Nov. 2002), 1344–1357. 1

[SWB*00] SATO Y., WESTIN C., BHALERAO A., NAKAJIMA S., SHIRAGA N., TAMURA S., KIKINIS R.: Tissue classification based on 3D local intensity structures for volume rendering. *IEEE Transactions on Visualization and Computer Graphics 6*, 2 (2000), 160–180. 2

[TPD06] TAPPENBECK A., PREIM B., DICKEN V.: Distance-based transfer function design: Specification methods and applications. In *Simulation und Visualisierung (SimVis)* (Magdeburg, Germany, 2006), SCS Publishing House e.V., pp. 259–274. 2

[VHHFG05] VEGA-HIGUERA F., HASTREITER P., FAHLBUSCH R., GREINER G.: High performance volume splatting for visualization of neurovascular data. In *IEEE Visualization, 2005. VIS 05* (Oct 2005), pp. 271–278. doi:10.1109/VISUAL.2005.1532805. 2