Eurographics/ IEEE-VGTC Symposium on Visualization (2006)
Thomas Ertl, Ken Joy, and Beatriz Santos (Editors)

The α-histogram: Using Spatial Coherence to Enhance Histograms and Transfer Function Design

Claes Lundström[†], Anders Ynnerman[‡], Patric Ljung[‡], Anders Persson[§] and Hans Knutsson[¶]

[†]Center for Medical Image science and Visualization (CMIV), Linköping University, and Sectra-Imtec AB, clalu@imv.liu.se
[‡]Division for Visual Information Technology and Applications, Linköping University, {andyn,plg}@itn.liu.se
[§]Center for Medical Image science and Visualization (CMIV), Linköping University, andpe@cmiv.liu.se
[¶]Medical Informatics Group, Linköping University, knutte@imt.liu.se

Abstract

The high complexity of Transfer Function (TF) design is a major obstacle to widespread routine use of Direct Volume Rendering, particularly in the case of medical imaging. Both manual and automatic TF design schemes would benefit greatly from a fast and simple method for detection of tissue value ranges. To this end, we introduce the α-histogram, an enhancement that amplifies ranges corresponding to spatially coherent materials. The properties of the α-histogram have been explored for synthetic data sets and then successfully used to detect vessels in 20 Magnetic Resonance angiographies, proving the potential of this approach as a fast and simple technique for histogram enhancement in general and for TF construction in particular.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation; I.3.6 [Computer Graphics]: Methodology and Techniques; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism;

1. Introduction

Direct Volume Rendering (DVR) is a well established and developed technique in scientific visualization. In medical imaging, DVR is growing from a peripheral tool to a central and crucial tool due to the rapidly increasing size of standard data sets making slice-by-slice viewing difficult [And03]. One of the major remaining challenges for DVR is the high complexity of Transfer Function (TF) design. The manual effort required is often prohibitive for routine use.

Histograms are widely used to characterize data sets, for example in TF design, thanks to their high information content combined with a simple form. In order to render a data set, the value ranges of the interesting features must be known. For uncalibrated capture techniques, such as Magnetic Resonance Imaging (MRI), this is not the case. Nevertheless, if the features would appear as clearly visible peaks in the histogram both manual and automatic TF construction would be simple. It often occurs, however, that the feature peaks cannot be distinguished in the histogram. One reason is that the background distribution may completely cover a minor feature, see figure 1, another reason is that two materials may have overlapping intensity ranges.

In this paper we present a novel way to incorporate spatial coherence into an enhanced histogram: the α-histogram. The data set is divided into local regions for which individual histograms are retrieved. These local histograms are then raised to the power of α before summation and normalization. Thus, value ranges with high spatial coherence, i.e. spatially concentrated, will be amplified in the resulting histogram, effectively enlarging peaks corresponding to distinct materials. In order to exploit the enhancements in the α-histogram we propose an automatic peak detection scheme based on an adaptive area measure.

The results show that the α-histogram is a fast and simple general-purpose histogram enhancement method, with potential use in any application that processes or displays a histogram. Furthermore, we have successfully used the α-histogram in a scheme for automatic detection of tissue value ranges in medical volumes, demonstrating higher accuracy and robustness than alternative methods.

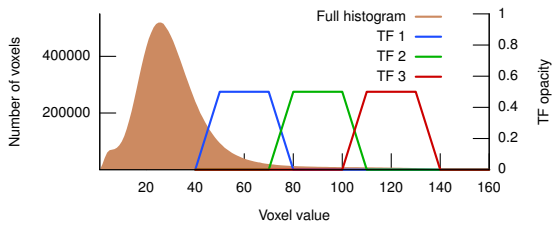

Figure 1: *Difficult TF construction. For this MR angiography there is no clue to the value range of the vessels in the full histogram. Thus, there is no guidance for selecting an appropriate TF, which in this case would be number 3.*

2. Related work

The challenge of simplifying TF construction has been addressed in many different approaches. One solution is to let the user explore the TF space by iteratively selecting the most appropriate visualization from a "gallery" of thumbnails [HHKP96]. Another approach is to base the TF on materials defined by implicit high-dimensional attributes [TLM03]. The attributes are derived by a neural network trained by example regions defined by the user.

The Contour Spectrum method [BPS97] provides an automatic analysis guiding the user to select appropriate values for isosurfacing. Semi-automatic TF generation for DVR of boundary regions has also been proposed [KD98]. Displaying only material boundaries is, however, not sufficient in medical visualization, primarily since tissue density is diagnostically important and secondly since the boundaries are often highly distorted by noise.

Many tissue classification schemes are focused on a particular examination type, e.g. MR brain images, while general-purpose methods are more relevant to discuss here. Nyúl and Udupa [NU99] proposed a standardization scheme for the intensity scale of MR images, implemented for histograms with a single main peak. Scale segments are linearly transformed according to landmarks retrieved from the histogram. Another standardization scheme for MR brain images analyzes higher-order derivatives [Chr03]. Lundström et al. [LLY05] introduced Partial Range Histograms (PRHs) for tissue detection purposes, combining intensity ranges with spatial regions. Their scheme divides the global histogram into separate tissue histograms and has been used for TF adaptation in a clinical evaluation [PBL*]. The α-histogram is an alternative to PRHs in some situations, and the two methods are compared in this paper.

Combinations of histograms with spatial relations, although not directly relevant to this paper, have been used in other research areas: content-based image retrieval [PZ96], registration [NUS03], segmentation [ZBS01], and sub-voxel classification [LFB98]. In an extension to 2D TFs for DVR, a grouping of histogram points based on spatial coherence has been proposed [RBS05]. A fundamental difference, both from α-histograms and PRHs, is that features with overlapping value ranges cannot be separated by this technique.

3. The α-histogram framework

The main limitation of any histogram is that all spatial relations in the data set are lost. The α-histogram, on the other hand, exposes a highly valuable type of spatial information, namely which value ranges correspond to spatially coherent features. This is valuable since the goal of a TF is typically to capture such features and suppress the other data. Thus, the histogram is enhanced while its simple form is retained.

3.1. α-histogram definition

We define a histogram incorporating spatial coherence, the α-histogram. It is constructed as a sum of local histogram functions $H_n(N,x)$, counting the instances of the value x in a local neighborhood (eq. 1). N is the set of data values in the neighborhood, D_x is the set of data values equalling x in the data set D. $|S|$ is the cardinality of set S.

$$H_n(N,x) = |N \cap D_x| \tag{1}$$

The first step is to divide the data set into spatial regions $N_1,\ldots,N_k$. We will define the α-histogram for a non-overlapping neighborhood subdivision:

$$D = \bigcup_{i=1}^{k} N_i \quad \text{and} \quad N_i \cap N_j = \emptyset, \quad i \neq j$$

If the region contains coherent tissues there will be small peaks for those value ranges in the local histogram. In contrast, the majority of regions containing only noise and incoherent material will not show these distinct peaks. The idea is now to amplify the local tissue peaks by raising the histogram values to the power of $\alpha > 1$. The α-histogram $H_\alpha(x)$ is the sum of all the enhanced local histograms (eq. 2) in the data set D. An important special case is to set $\alpha = \infty$, this is equal to a maximum operator. An initial normalization is made by raising the summed histogram values to $1/\alpha$.

$$H_\alpha(x) = \left(\sum_{i=1}^{k} H_n(N_i,x)^\alpha \right)^{1/\alpha} \tag{2}$$

It is convenient to introduce a second normalization, making the α-histogram have the same area as the original histogram (eq. 3). In this way, displaying the histogram for increasing α will result in growing peaks for coherent tissues at the expense of shrinking peaks of low importance. The benefits are that the scale of the histogram display can be fixed and that the change is easy to perceive visually.

$$\widetilde{H}_\alpha(x) = \frac{|D|}{\sum_x H_\alpha(x)} \; H_\alpha(x) \tag{3}$$

In general terms, the α-histogram is a modulated summation of local histograms. As such, it relates to other summation

techniques. The underlying idea is the same as for Root-Mean-Square-Error (RMSE): outlier values get a higher impact by squaring the values before summing them. The final square root is a normalization, bringing back the original scale. In RMSE averaging is applied, in contrast to the α-histogram case where the sum is the wanted attribute.

An image quality measure is often a sum of partial errors from different spatial regions and frequency bands. The traditional summation approach is known as *Minkowski pooling* [WBSS04] (eq. 4, where l is the spatial index and m denotes the frequency band). The individual errors are raised to a factor β before summation, and the result is normalized by taking the β root. Hence, for each value x, the α-histogram corresponds to a Minkowski pooling of the local histograms with $\beta = \alpha$.

$$E = \left(\sum_l \sum_m |e_{l,m}|^\beta \right)^{1/\beta} \tag{4}$$

An important note is that the α-histogram makes no assumptions about the type of the data set. This paper mainly concerns data sets with a structured three-dimensional grid, but the framework is generally applicable.

3.2. Variations

There are a number of possible variations of the standard α-histogram that can exploit *a priori* information on the data set and the task at hand. Spatial coherence is represented by the local histograms. At this stage any type of histogram features can be enhanced through filtering. A typical task is to find tissues at the high end of the intensity scale, which can be achieved by selectively amplifying these ranges. A similar effect can be achieved by increasing α for increasing x. If border regions are of interest then local histograms with more than one peak can be selectively amplified.

The type of neighborhoods to use for the local histograms has a clear impact on the result, the standard solution with non-overlapping cubical blocks is not mandatory. A feasible extension would be to use overlapping regions with a spatial weighting window that lets voxels near the region centre have larger impact on the local histogram. The shape of the neighborhood also matters. If the sought-for tissues are known to have a typical shape, an anisotropic neighborhood of the same shape will be a better detector than an isotropic one. The shape can also be automatically adapted to the local characteristics of the data. A brute-force approach is to sum amplified local histograms for anisotropic neighborhoods of 'all' directions, the neighborhood coinciding with the orientation of a spatially coherent feature would still dominate if α is large enough.

4. Peak detection

The α-histogram can be used simply as a guide in manual TF design. The main goal of this paper is, however, to automatically adapt TFs between data sets. If the typical values of the interesting materials are known, the TF can be related to these landmarks and easily be transformed for each new data set. Therefore, an automatic peak detection scheme for histograms is needed. The first challenge is that the α-histogram is usually quite noisy. When smoothing the curve care must be taken to only remove the unwanted irregularities without affecting the real peaks, which may be very subtle. Avoiding overly strong smoothing, an abundant number of peaks will remain, requiring an assessment to sort out the relevant ones. This assessment is not trivial, as seen in figure 2.

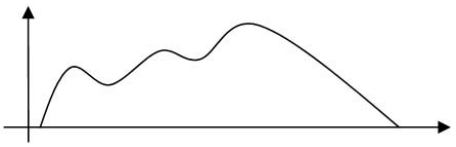

Figure 2: *Complex peak shapes. If this histogram is to be represented as two peaks, which would they be?*

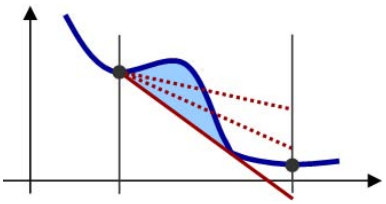

Figure 3: *Definition of peak area. The base line is the solid straight line, invalid base lines are dashed. The peak area is shaded.*

4.1. Peak importance

We define a peak as a histogram segment consisting of an apex surrounded by two valleys. An apex is a point with greater height than both its neighbors, the opposite is true for a valley. It is intuitive to connect the importance of a peak to its area. We choose to relate the area to an implicit base line, oriented according to the peak shape. We define the base line as starting at the highest valley with a downward slope and touching exactly one more point on the histogram curve within the peak range, see figure 3. The area is calculated by summing the height difference between the base line and the histogram between the two touching points.

4.2. Peak analysis scheme

Our histogram peak analysis consists of three steps. First, smoothing is selectively applied to each value representing a minimal peak or crease (the middle of three directly adjacent valley-apex-valley or apex-valley-apex points), repeated until no such cases remain. Second, the whole histogram is subject to repeated smoothing. The stopping criterion is given as

a maximum number of remaining peaks: n_{p1}, set to 20 in this paper. The first selective smoothing step reduces the need of general smoothing, in order to avoid removal of interesting, subtle peaks. The smoothing operator is in both cases a simple binomial kernel: $0.25 \cdot [1 \;\; 2 \;\; 1]$.

Finally, the remaining peaks are listed and a simplification is performed, reducing them to a predefined number n_{p2}. The peak with the least area (see above) is removed from the list. A peak neighboring the removed one may be extended to cover the removed peak if its area would increase. If both neighbors would gain from an extension then the one having the largest peak height is selected. In summary, the number of peaks is decreased, keeping the most significant peaks.

4.3. Detection evaluation

In order to evaluate the results of the peak detection, accuracy and confidence measures are needed. We will use the notations $\{x_-, x_0, x_+\}$ to describe a true peak, where x_0 is the position of the peak apex and the peak width is described by x_- and x_+, defined as the 10th and 90th percentiles of the peak histogram area. A detected peak, i.e. a peak approximation, is denoted $\{x_-^*, x_0^*, x_+^*\}$, corresponding to the valley-apex-valley values from the peak analysis.

Accuracy is measured through the peak precision error e_p, defined as the absolute difference between x_0^* and x_0, relative to the true peak width (eq. 5). For confidence, the peak area is a possible measure, but there are cases when it would be misleading. An example is a wide, flat peak that would have a large area but still provide a very uncertain prediction. Instead we define peak confidence, c_p, to be the relative height difference of the peak apex $H(x_0^*)$ and the highest of the two valleys $H(x_-^*)$ and $H(x_+^*)$, where $H(x)$ denotes the histogram function (eq. 6).

$$e_p = \frac{|x_0^* - x_0|}{x_+ - x_-} \tag{5}$$

$$c_p = \frac{H(x_0^*) - \max\left(H(x_-^*),\; H(x_+^*)\right)}{H(x_0^*)} \tag{6}$$

5. Results

The test results are divided into two parts. First, synthetic data sets are used to explore the properties of the α-histogram and its variations. These experiences are then applied to a clinical TF adaptation problem, where tissue detection is the key component. The standard α-histogram based on blocks with 8^3 voxels is used, unless stated otherwise.

5.1. α-histogram properties

Initially, tests have been performed on two synthetic data sets with controllable characteristics. The first data set is a tube spiralling around a torus, intended to imitate a vessel,

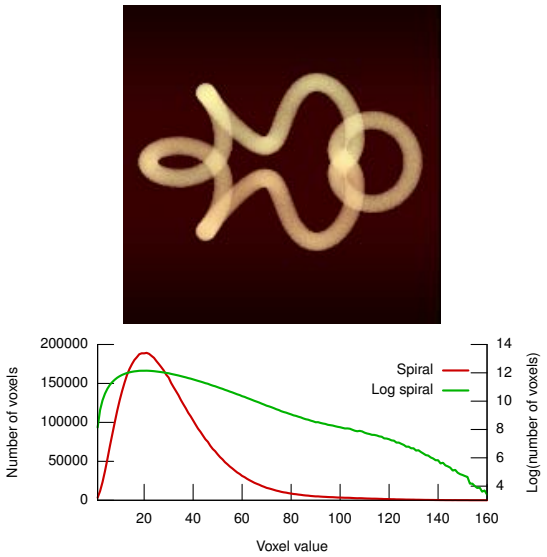

Figure 4: *Synthetic data set imitating a vessel, designed not to coincide with a Cartesian block structure. The background and the spiral have Gamma and Gaussian distributions, respectively. The spiral peak at value 100 is not visible, even using a logarithmic scale as proposed in [PM04].*

see figure 4. Since a block subdivision is performed in the standard α-histogram, the spiral shape is designed not to be aligned with any Cartesian grid. The radius, ρ, sets the width of the 'vessel'. The distributions of voxel values have been chosen to mimic real angiography data sets. The vessel has a Gaussian distribution ($\mu = 100$, $\sigma = 20$). The background has a Gamma distribution ($\mu = 30$, $\sigma = 10\sqrt{3}$) resulting in a single peak with no negative values but with large spread to high positive values.

The second synthetic data set imitates two organs, see figure 5. The curved 'yin-yang' shape is selected, as before, to be independent of the Cartesian blocking. The background has value zero and the two 'organs' both have Gaussian distributions ($\sigma = 30$), the first is fixed at $\mu_{y1} = 100$, whereas μ_{y2} can be varied. Both the synthetic data sets have 192^3 voxels, reconstructed for each test by assigning random voxel values according to the given distributions. Presented measures are averages over 10 trials, whereas histograms are from a single test. For these simulations a single measure to describe the quality of the α-histograms is desirable. We chose to use the peak confidence c_p with the added restraint that it is set to zero for peak detection failures (an apex offset larger than half a standard deviation of the true peak, or a valley being directly adjacent to the apex).

A first round of tests explored the effect of α on the peak detection ($n_{p2} = 2$). As seen in figure 6, the spiral peak becomes more apparent as higher α is applied. The disadvantage of higher α is the amplified noise. This noise makes the

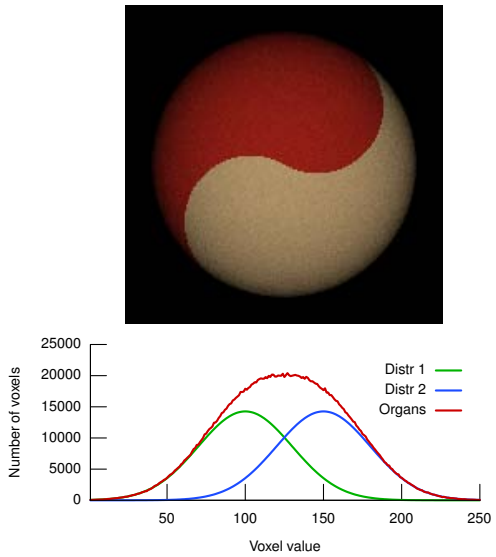

Figure 5: *Synthetic data set imitating two organs. The 'yin-yang' shape is chosen so as not to coincide with a Cartesian block structure. The background is zero, the 'organs' have separate Gaussian distributions.*

peak detection confidence drop for high α. A similar experiment was carried out on the synthetic organs, see figure 7. Peak detection confidence increases with higher α but the gain is most pronounced for difficult detection cases.

Apart from enabling peak detection, the α-histogram also enhances the visual appearance. In the spiral data set, as in many real data sets, the height of the background peak is orders of magnitude larger than the others, hindering a good overall view. Using, for example, logarithmic scale [PM04] can help but the α-histogram has the effect of both evening out the peak heights *and* promoting spatial coherence.

An important parameter in the block-based α-histogram is the block size. In figure 8 different block sizes are applied to spirals of different widths. Two conclusions can be drawn from the results: primarily that the size should not be smaller than 6^3 in order to obtain good effect from the spatial coherence and, secondly, that features of the same scale as the block size are favored by the α-histogram.

If there is *a priori* information that the tissue has larger extent in a certain dimension, the use of non-cubical blocks can be an advantage. This has been tested using a spiral data set where the spiral extends four times further in the z-direction than in x and y. Results from anisotropic blocks being 3-4 times larger in the z-direction were compared with those from cubical blocks. The results, see figure 9, show that there is some advantage to adapting the block shape to the tissue characteristics, but the effect is not dramatic.

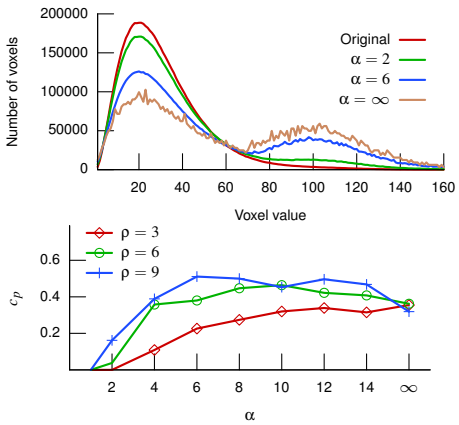

Figure 6: *Varying α, spiral data set. Top: For higher α the peak amplification and the noise both increase. (ρ = 6) Bottom: 'Vessel' peak detection. Detection improves for higher α but noise reduces the confidence at very high α.*

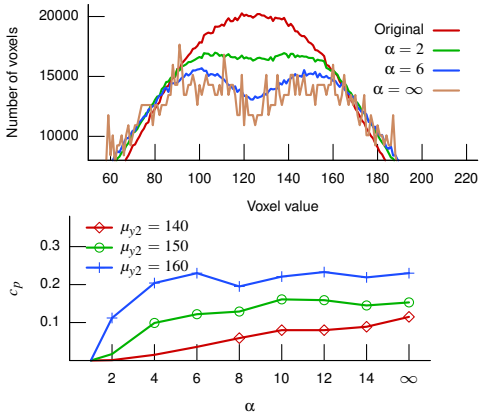

Figure 7: *Varying α, yin-yang data set. Top: For higher α the peak amplification and the noise both increase. ($\mu_{y2} = 150$) Bottom: Peak detection for the high-intensity 'organ'. Detection confidence generally increases with α but quickly reaches a plateau for simpler detections.*

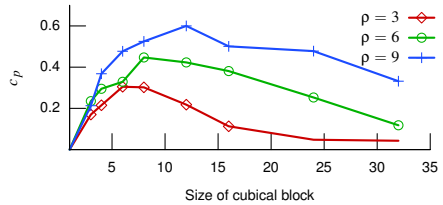

Figure 8: *Varying block size for spiral data sets, α = 10. Blocks smaller than 6^3 yield less effect from spatial coherence. Otherwise a block size similar to the size of the feature in question is appropriate.*

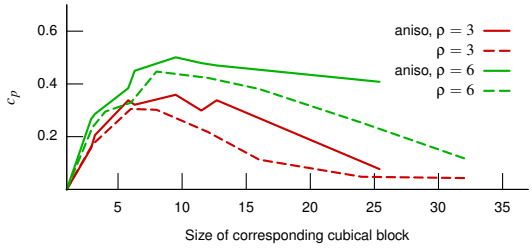

Figure 9: *Comparing anisotropic blocks to cubical blocks for stretched versions of the spiral data set ($\alpha = 10$). A block shape adapted to tissue properties yields better detection.*

Table 1: *Peak detection results for 20 MR AAA angiographies. The block size is 8^3 for both α-histograms and PRHs.*

		$\bar{e}_p$	$\tilde{e}_p$	$\bar{c}_p$
Percentile		0.21	0.17	-
Original histo		0.23	0.09	0.04
α-histo	$\alpha = 4$	0.11	0.02	0.19
	$\alpha = 10$	0.09	0.03	0.21
	$\alpha = \infty$	0.12	0.04	0.21
PRH	$\varepsilon = 0.6$	0.07	0.01	(0.76)
	$\varepsilon = 0.7$	0.14	0.08	(0.63)
	$\varepsilon = 0.8$	0.12	0.03	(0.73)

5.2. TF adaptation

The α-histogram has been applied to TF adaptation for a number of problematic clinical data sets. Primarily, 20 MR angiographies of Abdominal Aortic Aneurysm (AAA) have been tested, captured with several different protocols resulting in highly varying data set characteristics. The interesting tissue is the contrast agent-filled vessels, for which the typical intensity value varies between 182 and 1393 across the data sets. AAA images are typically used to prepare for surgical procedures where stents replace parts of the aorta. It is vital for the success of the operation to know the exact aorta diameter, an error of over 2 mm is usually considered clinically unacceptable. Inadequate TF parameters can cause diameter measurement errors well above this limit [PDE*04].

It is difficult to use the rendered images for a detailed quantitative evaluation of the TF adaptation, the gold standard available for these images is not precise enough. Instead we opt to compare the detected peak directly to the histogram of a manual segmentation of the vessels. As one of the alternative methods (PRH) has been successfully clinically evaluated [PBL*], we believe that its peak detection results set a good benchmark for the comparisons in this paper. The ground truth segmentations were validated by an experienced radiologist.

As a first reference, we have used a static peak apex predictor defined as the value corresponding to 98.1% of the full histogram area. The true apex in average corresponds to this percentile in these data sets. Intensity value 0 is discarded, as it is often orders of magnitude more frequent than the other values. Both the original histogram and a number of α-histograms have been evaluated using the peak detection scheme from section 4.2. As the appropriate choice of n_{p2} may vary, we applied four different simplifications with $n_{p2} = 1,\ldots,4$, respectively. Among this total of 10 peak definitions, those having the apex within a valid range were considered and the one with highest confidence was selected. If no peak was found within the range, the one closest to the range was used. The valid range was defined as the points corresponding to 96.5% and 99.2% of the histogram area, which is the range of the true apices in these data sets.

Finally, we have employed the Partial Range Histogram (PRH) method of our previous work, we refer to [LLY05] for a full description. The aim of the PRH method is to create a separate histogram for each tissue, in contrast to the α-histogram that makes tissues stand out at the global histogram level. In spite of the different approaches, both methods can be used for peak detection. The main parameter of the PRH method is the range weight limit, ε. The resulting number of PRH peaks, n_{PRH}, is typically between 5 and 10 and can be simplified further. In these data sets the vessel peak is known to be at the top of the intensity scale. We create a number of possible top peaks by merging $1,\ldots,n_{\mathrm{PRH}}-1$ of the topmost peaks. Then peak selection is performed according to the rules connected to the valid range described above.

The overall results of the peak detection are given in table 1. Accuracy is represented by precision error, both mean ($\bar{e}_p$) and median ($\tilde{e}_p$), and the robustness by the mean peak confidence ($\bar{c}_p$). Compared to the percentile reference, peak detection on the original histogram has lower error for a majority of cases, but the average error remains high.

The α-histograms provide clear improvement on peak detection accuracy as well as robustness. The results are consistent across all three values of α. The resulting α-histograms for two of the most difficult cases are presented in figure 10 and demonstrate that even if the peak apex detection is imprecise, the α-histogram can still be of great value in emphasizing the tissue value range. Some examples of TF adaptation based on α-histograms are given in figure 11.

As seen in table 1, the peak detection based on PRHs has the same level of quality as the α-histograms. Note that the peak confidence c_p is of little interest for PRHs; it is always very large since PRHs correspond to separated tissue peaks. For this peak detection task the differences between the PRH and α-histogram are minor, but in other situations the benefits of each method are more apparent. The PRH method makes a number of distinct predictions during the tissue detection process regarding position and width of histogram

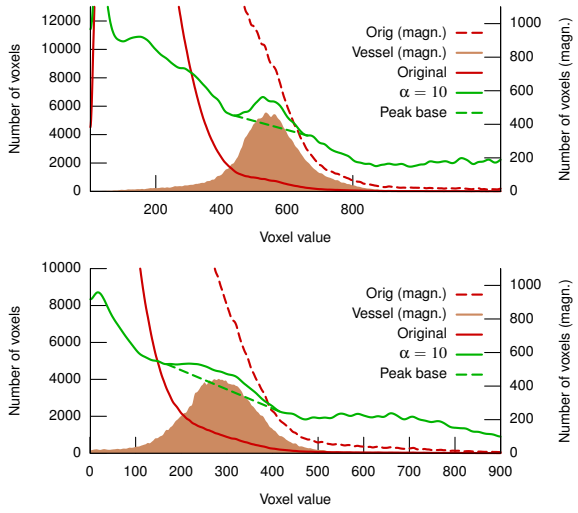

Figure 10: *Resulting α-histograms for two challenging MR angiography cases. Smoothing applied, $\alpha = 10$. Top: Using the α-histogram the peak is found with high accuracy and confidence ($e_p = 0.03$, $c_p = 0.20$). Bottom: Even though the detected peak apex is not quite accurate, the peak range is readily found ($e_p = 0.30$, $c_p = 0.01$).*

peaks. A slight error in such a prediction can propagate and grow into a major error in the final result. An indication of this weakness is the sensitivity to the choice of ε, in table 1 manifested as greatly differing $\widetilde{e}_p$. In contrast, the α-histogram is a robust general enhancement that defers the concrete tissue detection to the user or other algorithms.

The major advantage of PRHs is that they separate the tissue peak from the global histogram, making PRHs well suited to detection of highly overlapping p eaks, as shown by the example given in figure 12. This is an MR biliary duct examination where three distinct peaks can be found: liver, kidneys, and spleen. The α-histogram fails to clearly make the kidney peak stand out, whereas the PRH method correctly finds the position of all three peaks.

The sizes of the clinical data sets were 6-32 MB and an α-histogram with $\alpha = 10$ and block size 8^3 was calculated in 0.4-2.2 s on a 1.8 GHz Pentium M laptop. The PRH method had similar performance: 0.4-2.9 s. These times are quite acceptable, since the methods are typically applied in a preprocessing step.

6. Conclusions

We have presented the α-histogram, a novel tool in the continuing effort to facilitate TF design in DVR. The α-histogram introduces an automatic enhancement of spatially coherent value ranges, hence emphasizing distinct materials in the data set. One efficient variant is based on non-

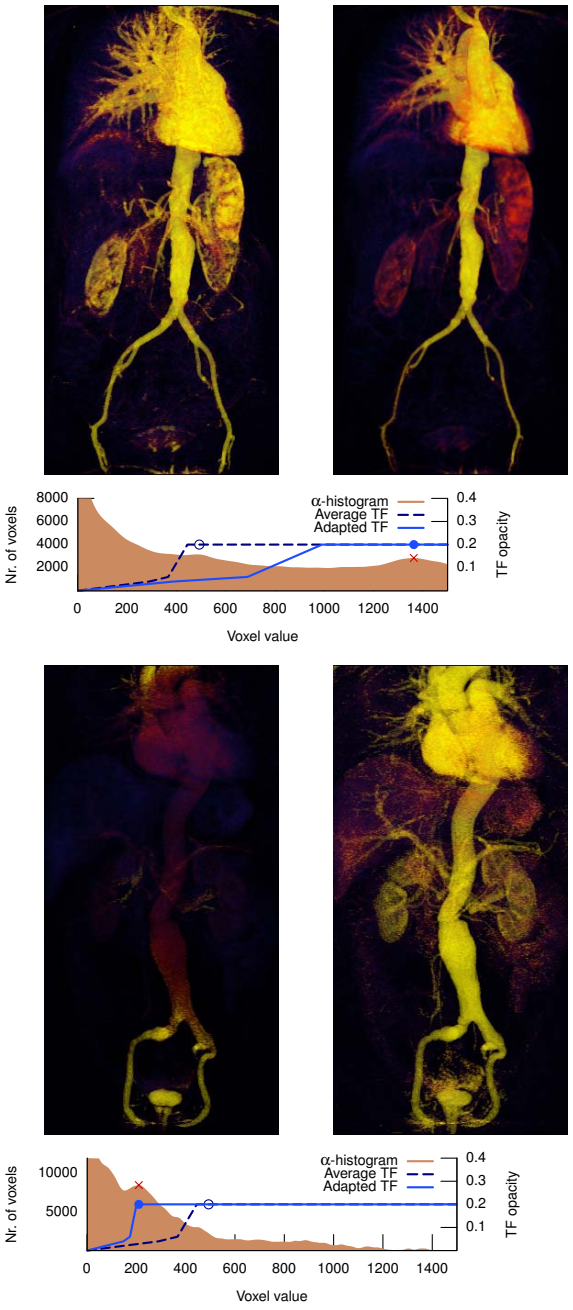

Figure 11: *Adapting TFs for MR angiographies. The width of the aorta (the large central vessel) is to be measured at high precision, the error should be less than 2 mm (≈ 2 pixels). Left column: A TF based on the average vessel intensity value of the 20 data sets is unacceptable. In the top image the aorta width is exaggerated, in the bottom image the aorta is hardly visible. Right column: An adapted TF based on the detected peak in the α-histogram ($\alpha = 10$) correctly renders the width of the aorta in both cases. In the graphs, a cross denotes the peak and a circle denotes the TF reference point.*

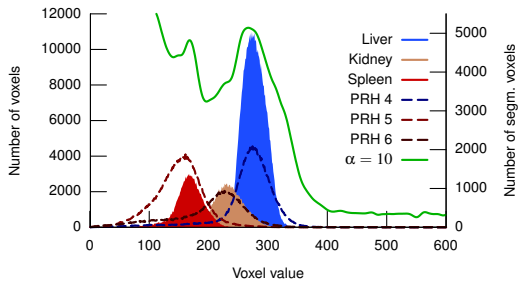

Figure 12: *Peak detection for MR biliary duct examination. In contrast to the PRH method, the α-histogram fails to distinguish the kidney peak, since it is highly overlapping.*

overlapping cubical block neighborhoods. Merely two parameters then exist, α and the block size, and the enhancement is not sensitive to the choice of these parameters.

The typical intensity values for vessels in MR angiographies have been detected automatically with high accuracy and confidence using the α-histogram. The enhancement can be used both to reveal invisible histogram peaks and also to increase the accuracy and confidence of distinguishable peaks. Our conclusion is that the α-histogram is appropriate as a general-purpose histogram enhancement, particularly useful for automatic TF adaptation between data sets.

In our future work, we will use the α-histogram framework, including the presented variations, to address other specific clinical visualization problems. The combination of α-histograms and PRHs is of particular interest.

Acknowledgements

This work has been conducted at the Center for Medical Image Science and Visualization (CMIV) at Linköping University, Sweden. CMIV is acknowledged for provision of financial support and access to leading edge research infrastructure. The work was funded by the Swedish Research Council, grants 621-2003-6582 and 621-2001-2778, and the Swedish Foundation for Strategic Research, grant A3 02:116. There are pending patent applications based on this work.

References

[And03] ANDRIOLE K. P.: A position paper from the SCAR TRIP(tm) subcommittee. http://www.scarnet.org/pdf/-TRIPwhitepaper1103.pdf, November 2003. Acquired March 2004.

[BPS97] BAJAJ C., PASCUCCI V., SCHIKORE D. R.: The Contour Spectrum. In *Proceedings IEEE Visualization 1997* (1997), pp. 167–173.

[Chr03] CHRISTENSEN J. D.: Normalization of brain magnetic resonance images using histogram even-order derivative analysis. *Magnetic Resonance Imaging 21*, 7 (September 2003), 817–820.

[HHKP96] HE T., HONG L., KAUFMAN A., PFISTER H.: Generation of transfer functions with stochastic search techniques. In *Proceedings IEEE Visualization 1996* (1996), pp. 227–234.

[KD98] KINDLMANN G., DURKIN J. W.: Semi-automatic generation of transfer functions for direct volume rendering. In *Proceedings IEEE Symposium on Volume Visualization* (1998), pp. 79–86.

[LFB98] LAIDLAW D., FLEISCHER K., BARR A.: Partial-volume bayesian classification of material mixtures in MR volume data using voxel histograms. *IEEE Transactions on Medical Imaging 17*, 1 (1998), 74–86.

[LLY05] LUNDSTRÖM C., LJUNG P., YNNERMAN A.: Extending and simplifying Transfer Function design in medical Volume Rendering using local histograms. In *Proceedings IEEE/EuroGraphics Symposium on Visualization* (2005), pp. 263–270.

[NU99] NYÚL L. G., UDUPA J. K.: On standardizing the MR image intensity scale. *Magnetic Resonance in Medicine 42* (1999), 1072–1081.

[NUS03] NYÚL L. G., UDUPA J. K., SAHA P. K.: Incorporating a measure of local scale in voxel-based 3-D image registration. *IEEE Transactions on Medical Imaging 22*, 2 (February 2003), 228–237.

[PBL*] PERSSON A., BRISMAR T., LUNDSTRÖM C., DAHLSTRÖM N., OTHBERG F., SMEDBY Ö.: Standardized volume rendering for MR angiography measurements in the abdominal aorta. *To appear in Acta Radiologica*.

[PDE*04] PERSSON A., DAHLSTRÖM N., ENGELLAU L., LARSSON E.-M., BRISMAR T., SMEDBY Ö.: Volume rendering compared with maximum intensity projection for magnetic resonance angiography measurements of the abdominal aorta. *Acta Radiologica 45* (2004), 453–459.

[PM04] POTTS S., MÖLLER T.: Transfer functions on a logarithmic scale for volume rendering. In *Proceedings of Graphics Interface 2004* (May 2004), pp. 57–63.

[PZ96] PASS G., ZABIH R.: Histogram refinement for content-based image retrieval. In *Proceedings 3rd IEEE Workshop on Applications on Computer Vision* (December 1996), pp. 96–102.

[RBS05] ROETTGER S., BAUER M., STAMMINGER M.: Spatialized transfer functions. In *Proceedings IEEE/EuroGraphics Symposium on Visualization* (2005), pp. 271–278.

[TLM03] TZENG F.-Y., LUM E. B., MA K.-L.: A novel interface for higher-dimensional classification of volume data. In *Proceedings IEEE Visualization 2003* (2003), pp. 505–512.

[WBSS04] WANG Z., BOVIK A. C., SHEIKH H. R., SIMONCELLI E. P.: Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing 13*, 4 (April 2004).

[ZBS01] ZHANG Y., BRADY M., SMITH S.: Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Transactions on Medical Imaging 20*, 1 (January 2001), 45–57.