Mirrored LH Histograms for the Visualization of Material Boundaries

Petr Šereda[1], Anna Vilanova[1] and Frans A. Gerritsen [1,2]

[1]Department of Biomedical Engineering, Technische Universiteit Eindhoven, Postbus 513, 5600 MB Eindhoven, The Netherlands Email: {P.Sereda,A.Vilanova}@tue.nl
[2]Advanced Development, Healthcare Informatics, Philips Medical Systems Nederland BV, P.O. Box. 10000, 5680 DA Best, The Netherlands Email: Frans.Gerritsen@philips.com

Abstract

The quality of direct volume rendering is strongly influenced by the transfer function. The purpose of the transfer function is to highlight aspects of the volume data that are of interest to the user. Material boundaries are often regarded as interesting, since they reveal the shapes of objects and their spatial relations.

We use a new transfer function domain that facilitates the selection of material boundaries in the volumetric data sets. Our domain enables to distinguish between both sides of the boundary. By reducing the partial volume effect on the boundary, we are able to easily assign optical properties to objects and their boundaries.

1 Introduction

The transfer function (TF) plays a crucial role in the direct volume rendering pipeline. The TF typically assigns optical properties such as color and opacity to every point of the data. A properly constructed TF will help to visualize those parts of the volumetric data set that are of interest to the user and will remove the rest by assigning a low opacity. By assigning different colors, the TF helps to show different objects in contrast.

For many applications, the attention of the user is focused on the boundaries of objects. The boundaries can reveal important information such as the shape of the object, its extent, size and spatial relation with other objects. In order to effectively select the data points that lie on the boundary and to differentiate between boundaries, one needs to use an appropriate TF domain. It has been shown in previous research [5, 6, 14] that boundaries cannot be well classified by using the scalar value only. Additional information such as the gradient magnitude or the boundary profile helps to differentiate boundaries.

Current TF domains enable the selection of a boundary and the assignment of certain color and opacity values. Every boundary between two materials, however, belongs partially to both materials. If the user wanted to visualize both materials by assigning different optical properties to their boundaries, the classification of the common boundary would be ambiguous.

We present a TF domain in which the boundaries are divided into two parts by the edge location. Our method extends the possibilities of the LH space [14] by enabling the selection of divided boundaries. In the LH space the divided boundaries are positioned in such a way that their relation is intuitive and the user can easily select all boundaries of certain material. We further show that by a simple projection of the extended LH space one can obtain a new 1D TF domain that further facilitates selection of materials including their boundaries.

2 Related Work

The transfer functions have been an extensively researched topic. We will focus on the related work that aims on the visualization of material boundaries. For a more detailed overview the reader is referred to, e.g., [4].

Levoy [9] introduced transfer functions for displaying surfaces in volumetric data. For an automatic selection of the important iso-contours Bajaj et al. [1] and Pekar et al. [12] used the contour spectra and Fujishiro et al.[3] investigated the behavior of the contours.

Due to the partial volume effect on the boundaries, one often cannot distinguish material boundaries using the scalar value only. Laidlaw et al. [8] used model-based Bayesian classification in order

to deal with partial volume voxels on the boundaries. Kindlmann and Durkin [5] showed that by combining the scalar value and the gradient magnitude one may distinguish between boundaries. In this 2D space the boundaries appear as arches. Kniss et al. [6, 7] used this domain for their 2D transfer functions. Although the arches enable some distinction between boundaries, they often overlap causing selection ambiguities. Lum et al. [10] used two additional samples on the boundary in order to improve on that. Šereda et al. [14] went further and traced the boundary profiles in every voxel to determine the start and end of the arches. These two intensity values are the coordinates of the voxel in the so-called LH space. It was shown that in the LH space the boundaries appear as blobs and are easier to separate then the arches. This property was further exploited [13] by using the shapes of the blobs in a clustering approach.

In this paper we exploit the advantages of the LH space and introduce its extension that we call the mirrored LH space. In this mirrored LH space both sides of the boundary can be classified independently. Furthermore, this space enables an easy identification of all boundaries that belong to certain material.

We also show that a projection of the mirrored LH histogram into 1D results in a histogram with resolved partial volume voxels on the boundaries. Therefore the peaks in this histogram appear to be more clear. The work of Lundström et al. [11] also aims on emphasizing histogram peaks, however, their approach is based on emphasizing locally coherent structures.

3 Mirrored LH Histogram

In this section we will first briefly describe the process of projecting the data samples into the LH space and construction of the LH histogram. We point out the most relevant properties of the LH space. Then we show the new extension of the LH space in which the LH histogram appears as two symmetric copies of the original histogram. For a more detailed description of the LH space we refer the reader to [14].

3.1 LH Histogram

For each point in the volume that lies on a boundary between two materials, the intensity values of the materials are determined. The lower intensity F_L and the higher intensity F_H can be found by examining the intensity profile across the boundary. The F_L and F_H are then the coordinates of the point in the 2D LH space. Figure 1 illustrates the boundary profile tracking approach that identifies both values.

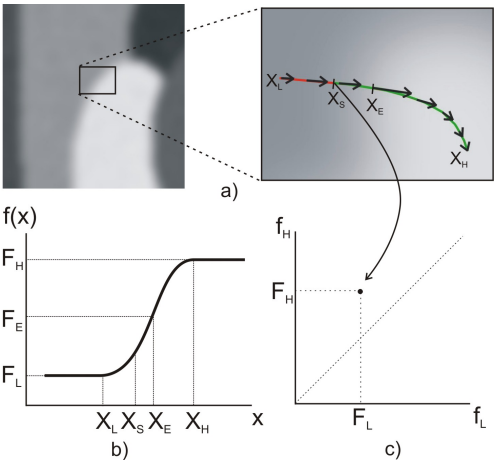

Figure 1: a) Starting at position X_S we generate a path across the boundary by integrating the gradient field. The positive gradient direction leads to X_H. Following the opposite direction ends by finding X_L. b) The intensity profile along the path. At points X_L and X_H where the tracking stops we read the values F_L and F_H respectively. F_E is the intensity value at the edge location X_E. c) F_L and F_H are used as coordinates in the LH space.

Figure 2 shows three examples of boundary profiles. The stopping criteria of the boundary tracking algorithm can handle any combination of these profiles. Points that do not lie on any boundary result in an early termination of the tracking and $F_L = F_H$.

All points lying on the same boundary profile project into one point. By projecting all voxels of the data into the LH space, one can obtain a two dimensional histogram called the LH histogram. Since by definition $F_H \geq F_L$, all boundary points project in the upper halve of the histogram. In the LH histogram, the boundaries and constant areas appear as blobs (local peaks). The width of the blobs depends on the variance of the intensity values present in the material. The height of the blobs (the amount of voxels contributing to the blob) is related to the surface of the boundaries.

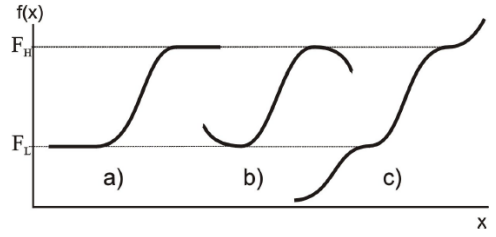

Figure 2: Three types of intensity profiles across the boundary. Boundary ends by a) large constant areas, b) local extrema or c) inflex points.

3.2 Division of the Boundary

The standard LH space is designed to classify the whole boundary at once. Every boundary does, however, exist between two materials and is therefore shared by both of them. Figure 4 illustrates the problem that originates in shared boundaries. If we want to visualize two neighboring materials, we need to decide how to classify their common boundary. Figure 4b shows the two critical boundaries between objects 1-4 and 4-7 (labeled by ovals). In this example the boundaries were assigned to the object 4. As the opacity of this object decreases, so does the opacity of the common boundaries. Therefore the objects 1 and 7 appear to have holes since part of their boundary is missing (Figure 4d).

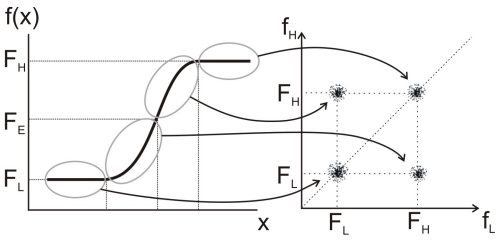

Figure 3: The voxels lying above the edge intensity project into the upper section with coordinates $[F_L;F_H]$. The voxels with lower intensity project below the diagonal at $[F_H;F_L]$.

In order to differentiate between both sides of the boundary, we divide the boundary profile into two parts by the edge. The edge location is determined during the boundary profile tracking by looking at the zero crossing of the second derivative in the gradient direction. There are, however, more advanced methods for locating edges, especially on highly curved surfaces [2].

Voxels lying above the edge, i. e., voxels with intensity $f(x) > F_E$, where F_E is the intensity at the edge location X_E (see Figure 1), project above the diagonal (same as in the standard LH space). Voxels lying below the edge, i.e., voxels with $f(x) > F_E$, project below the diagonal. That is achieved by swapping their F_L and F_H so that $F_H \leq F_L$ (see Figure 3).

Since there is approximately the same number of voxels on both sides of the boundaries, the lower part of the extended LH histogram seems to be a mirrored version of the upper part.

3.3 Properties

The phantom data set in Figure 5 illustrates the possibilities of the mirrored LH space. Both parts of the split boundary can be selected independently. That allows to select only the part of the boundary that belongs to the object and resolve the ambiguous selections. Especially interesting is the fact that all partial boundaries belonging to the same material are horizontally aligned in the mirrored LH space. That allows us to easily select all relevant boundaries (see Figure 5b).

The boundary division allows for visualization of boundary voxels lying on one side of the edge. Therefore, assuming the edge is an appropriate estimate of objects' borders, our new approach may more accurately depict the true size and shape of the objects.

3.4 Horizontal Projection

As shown above, the constant regions inside the material and the boundaries of the same material are horizontally aligned in the LH space. Therefore one could, in principle, select all voxels belonging to a certain material knowing only F_H. By a horizontal projection onto the F_H axis, we can discard the F_L values and generate a 1D histogram. In the new projected histogram one can still select all boundaries belonging to one material although the second material that forms the boundary cannot be identified anymore. As shown in Figure 6 the user could then make a simple selection in the 1D histogram, which is equivalent to a horizontal selection in the mirrored LH histogram.

The clarity of the peaks in standard 1D histograms is often hampered by the partial volume

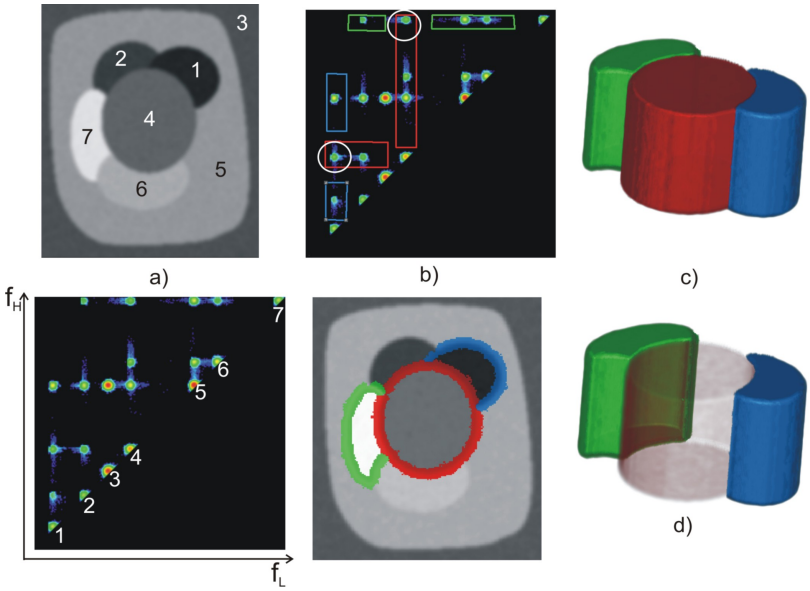

Figure 4: A phantom data set containing seven materials. a) A slice of the data set and the standard LH histograms with labeled interiors of materials. The constant areas project as blobs onto the diagonal. The other blobs correspond to boundaries. b) Boundaries of materials 1, 4 and 7 are selected. The classification of the two boundaries (1-4 and 4-7) labeled by an oval is ambiguous. c) and d) show rendering with full and partial opacity of the cylindrical object 4.

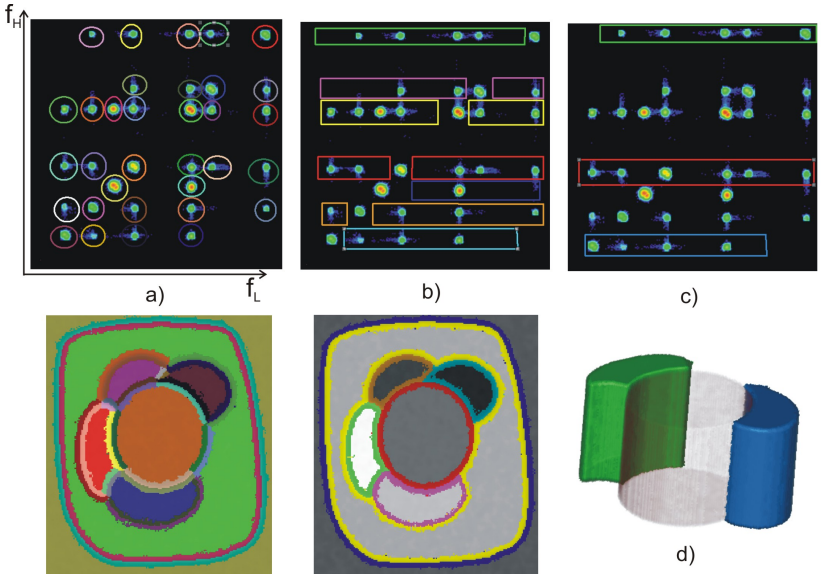

Figure 5: The possibilities of the mirrored LHspace illustrated on a phantom. a) Partial boundaries can be selected independently. b) All (parts of) boundaries that belong to one material can be easily selected, since they lie on a horizontal line. For the sake of clarity the interiors have not been selected. c) Improving the selection from Figure 4. Three horizontal selections contain everything that belongs to each of the materials. d) The rendering with the cylindrical object set to semitransparent.

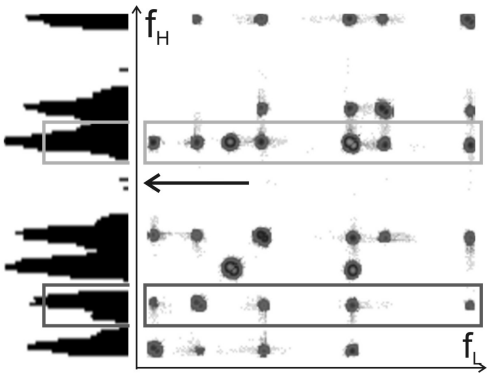

Figure 6: Discarding F_L information by horizontally projecting the LH histogram. The selections in both histograms are equivalent.

voxels on the boundaries. The mirrored and projected histogram, however, resolves these boundary voxels. As result, the peaks of the histogram become more clear. Figure 7 illustrates the mechanism of resolving the partial volume voxels on the material boundaries. Figure 8 shows the comparison of both histograms for the engine data set.

4 Results and Discussion

We will demonstrate the use of the mirrored LH space on several data sets. Figure 9 shows a rendering of the engine data set. On the left a selection in the standard LH space is used. The hard material-air and the hard-soft material boundaries are shown as part of the hard material components. On the right, however, the boundaries are divided and only the part of the boundary that belongs to the material is selected. It is clearly visible that the hard components appear to be wider in the left rendering, which might be misleading. In the right image only voxels that lie within the object boundaries (detected by the edge location) contribute to the rendering. Therefore the selection of boundaries in the mirrored LH space has the potential to better visualize the size and shape of objects.

Figure 10 shows the selection and rendering of three materials including their boundaries. When the soft material is set as semi-transparent one can see that the boundaries consist of two parts.

The tooth data set in Figure 11 is an example where a combined boundary selection can be used.

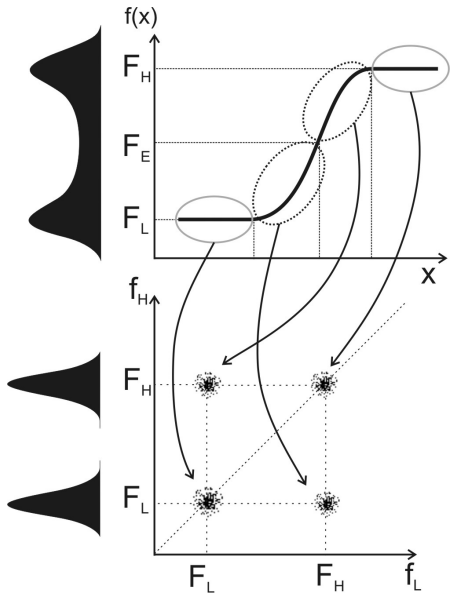

Figure 7: Resolving the partial volume voxels on the boundary. In the top the boundary contains scalar values in the range between F_L and F_H. These partial volume values (labeled by the dotted ovals) weaken the contrast of the material peaks in the histogram. In the mirrored LH space they are, however, mapped to F_L or F_H and therefore in the projected histogram emphasize the pure material peaks.

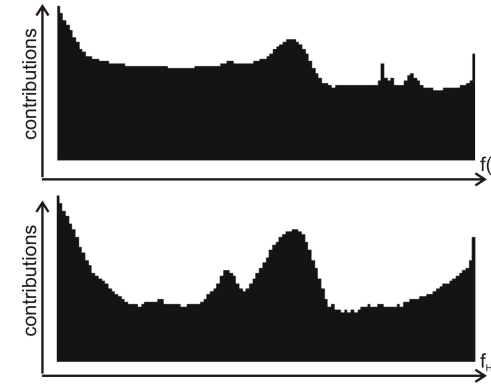

Figure 8: The engine data set. A comparison of the original histogram of scalar values and the projected histogram containing F_H values. By resolving the partial volume voxels on the boundaries, the peaks become more clear. The histogram contributions are shown in logarithmic scale.

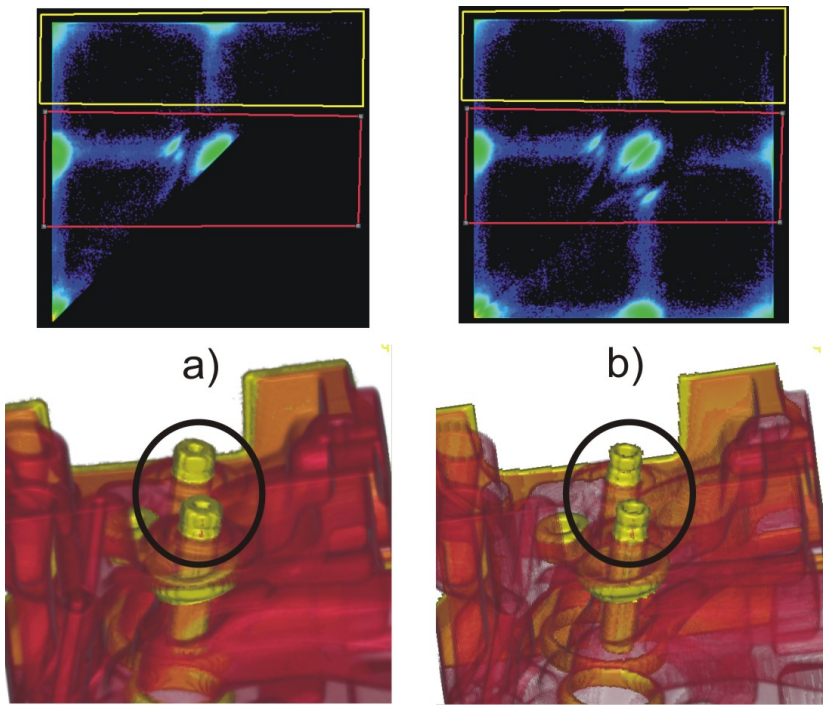

Figure 9: The engine data set. a) A TF based on the standard LH histogram. b) Selection of all voxels belonging to the same material in the mirrored LH histogram. In a) the objects appear thicker due to the contributions of voxels on both sides of the edge. This effect is well visible on the hard material components.

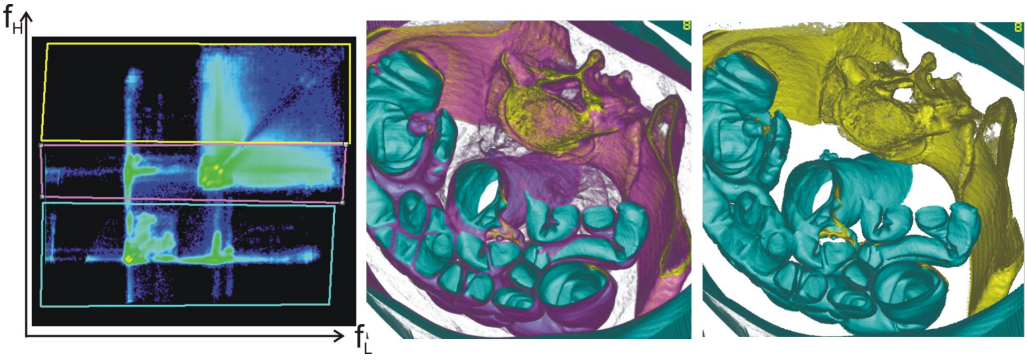

Figure 10: Abdominal CT scan. Three materials are highlighted: the air in the colon, the bone and the soft tissues. Notice the effect of both sides of the boundaries being classified separately. The right rendering shows only the colon wall and the pelvis bone.

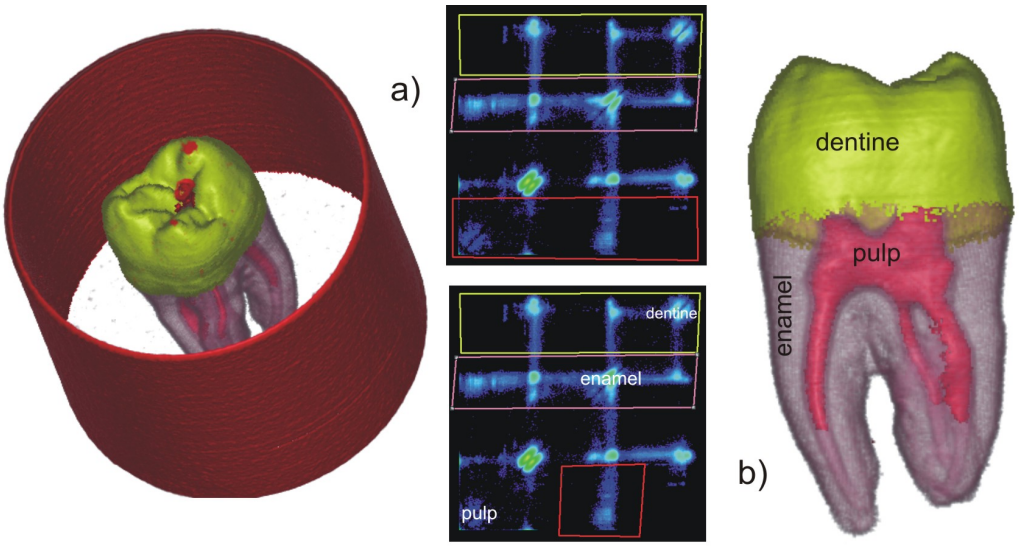

Figure 11: The CT tooth data set. The dentin and enamel boundaries are selected by using the horizontal selection of all relevant boundaries. In a) the pulp was selected by using the same technique. Since the pulp and background have similar intensities, in b) the selection is restricted only to the pulp-enamel boundary.

Two tissues are selected using the horizontal selection technique. The third tissue is selected by restricting the F_L values in order to avoid the selection of boundaries with the background.

Use of the mirrored LH space does not involve any additional performance or memory requirements as compared to the standard LH space. The only difference is that our tracking algorithm differentiates voxels lying above and below the edge. For those lying below the edge it returns F_L instead of F_H and vice versa. The rest of the visualization pipeline remains unchanged.

The hierarchical clustering method that has been shown in the previous work [13] can also be used in the mirrored LH space. Figure 12 shows one of the possible clustering of the mirrored LH histogram. The similarity measure based on the shapes of the histogram peaks was used. The clustering could, e.g., help to select boundaries of one material by automatically selecting all clusters lying on a horizontal line.

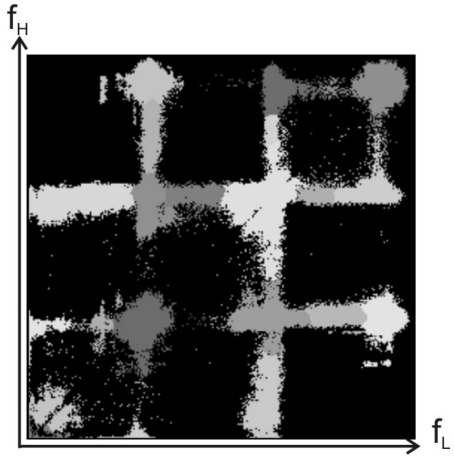

Figure 12: One of the possible clusterings of the mirrored LH histogram of the tooth data set.

5 Conclusions and Future work

We introduced a new way to classify material boundaries by using the mirrored LH space. Our

new approach helps to solve some of the classification ambiguities present in existing methods. The intuitive arrangement of boundaries in the mirrored LH space proved to facilitate the selection of all relevant boundaries. Furthermore, dividing the boundary by using the edge location prevents the voxels that lie outside object borders from being rendered. That helps to improve the visualization of the sizes and shapes of objects.

The projection of the mirrored LH space results in a 1D histogram that is not hampered by the partial volume voxels on material boundaries. This projection can help to simplify the TF by reducing its dimensionality. The 1D TF based on the projected F_H values can be used to make selections equivalent to the horizontal selections in 2D. An addition would be to offer the second dimension to the user only when needed.

The hierarchical clustering approach used in previous work could be adapted to the mirrored LH space. It could, e.g., preferably group clusters in the horizontal direction or consider similar grouping in both halves of the space.

Acknowledgment

The work presented in this paper has been funded by Philips Medical Systems in Best, The Netherlands.

References

[1] C. L. Bajaj, V. Pascucci, and D. Schikore. The contour spectrum. *Proceedings IEEE Visualization*, pages 167–174, 1997.

[2] H. Bouma, A. Vilanova, L. J. van Vliet, and F. A. Gerritsen. Correction for the dislocation of curved surfaces caused by the PSF in 2D and 3D CT images. *IEEE Transactions of Pattern Analysis and Machine Intelligence*, 27(9):1501–1507, 2005.

[3] I. Fujishiro, T. Azuma, and Y. Takeshima. Automating transfer function design for comprehensible volume rendering based on 3D field topology analysis. *Proceedings IEEE Visualization*, pages 467–470, 1999.

[4] A. Kaufman and K. Mueller. Overview of volume rendering. *chapter in The Visualization Handbook, Elsevier*, 2005.

[5] G. Kindlmann and J. W. Durkin. Semi-automatic generation of transfer functions for direct volume rendering. *Proceedings IEEE Symposium on Volume Visualization*, pages 79–86, 1998.

[6] J. Kniss, G. Kindlmann, and C. Hansen. Interactive volume rendering using multi-dimensional transfer functions and direct manipulation widgets. *Proceedings IEEE Visualization*, pages 255–262, 2001.

[7] J. Kniss, G. Kindlmann, and C. Hansen. Multidimensional transfer functions for interactive volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 8(3):270–285, 2002.

[8] D. H. Laidlaw, K. W. Fleischer, and A. H. Barr. Partial-volume bayesian classification of material mixtures in MR volume data using voxel histograms. *IEEE Transactions on Medical Imaging*, 17(1):74–86, 1998.

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

[10] E. B. Lum and K. L. Ma. Lighting transfer functions using gradient aligned sampling. *Proceedings IEEE Visualization*, pages 289–296, 2004.

[11] C. Lundström, A. Ynnerman, P. Ljung, A. Persson, and H. Knutsson. The α-histogram: Using spatial coherence to enhance histograms and transfer function design. *Proceedings IEEE/EuroGraphics Symposium on Visualization*, pages 227–234, 2006.

[12] V. Pekar, R. Wiemker, and D. Hempel. Fast detection of meaningful isosurfaces for volume data visualization. *Proceedings IEEE Visualization*, pages 223–230, 2001.

[13] P. Šereda, A. Vilanova, and F. A. Gerritsen. Automating transfer function design for volume rendering using hierarchical clustering of material boundaries. *Proceedings IEEE/EuroGraphics Symposium on Visualization*, pages 243–250, 2006.

[14] P. Šereda, A. Vilanova, I. W. O. Serlie, and F. A. Gerritsen. Visualization of boundaries in volumetric datasets using LH histograms. *IEEE Transactions on Visualization and Computer Graphics*, 12(2):208–218, 2006.

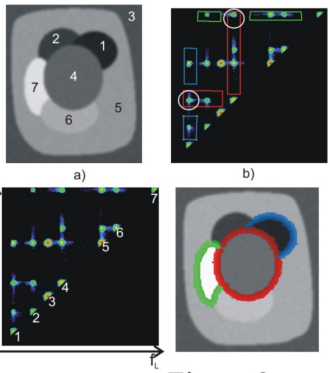

Figure 3

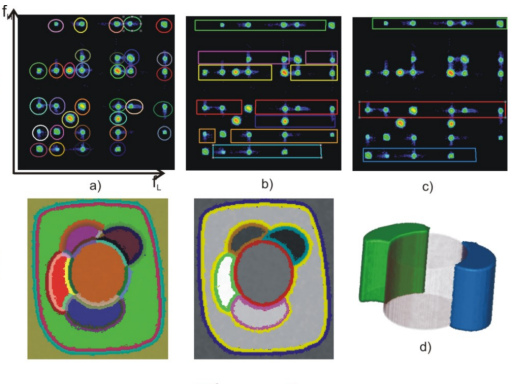

Figure 5

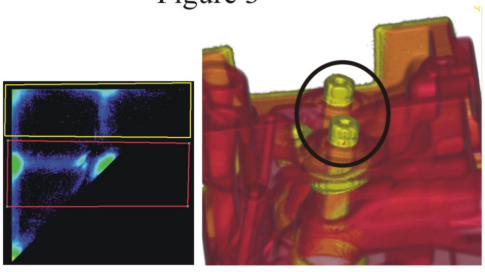

Figure 9

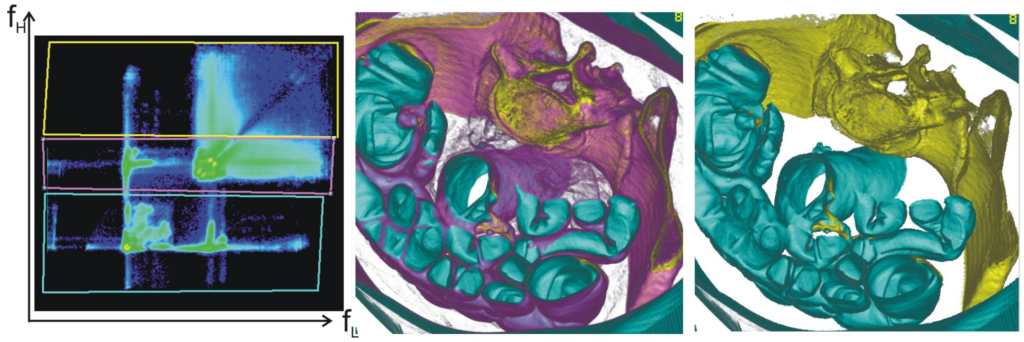

Figure 10

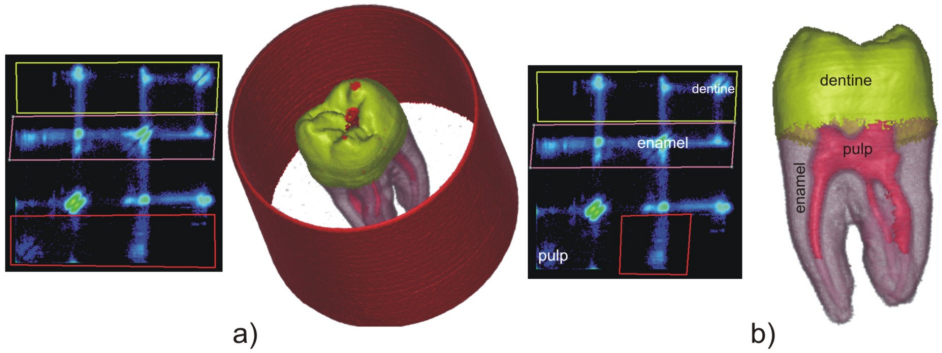

Figure 11