Moment Curves

Daniel Patel*
Christian Michelsen Research, Bergen, Norway

Martin Haidacher†
Institute of Computer Graphics and Algorithms, Vienna University of Technology, Austria

Jean-Paul Balabanian ‡
Department of Informatics, University of Bergen, Norway

Eduard M. Gröller §
Institute of Computer Graphics and Algorithms, Vienna University of Technology, Austria

ABSTRACT

We define a transfer function based on the first and second statistical moments. We consider the evolution of the mean and variance with respect to a growing neighborhood around a voxel. This evolution defines a curve in 3D for which we identify important trends and project it back to 2D. The resulting 2D projection can be brushed for easy and robust classification of materials and material borders. The transfer function is applied to both CT and MR data.

Index Terms: I.4.10 [Image Processing]: Image Representation—Volumetric, I.4.10 [Image Processing]: Image Representation—Statistical, I.4.6 [Image Processing]: Segmentation—Pixel classification, I.4.7 [Image Processing]: Feature Measurement—Feature representation

1 INTRODUCTION

We present a novel transfer function specification that uses the first and second statistical moments - the mean and variance. With these basic measures we are able to classify materials based on nontrivial properties, such as a material's internal variance. The transfer function is also robust to uniformly distributed noise and both identifies material boundaries and differentiates between them. The robustness of the method comes from the fact that it simultaneously considers the data on multiple scales when calculating the first and second statistical moments.

2 RELATED WORK

There have been many publications on transfer functions and on simplifying their design. Early publications focus on 1D transfer functions where color and opacity are assigned directly to the intensity values of the volume. The design galleries by Marks et al. [8] attempt to simplify the design of transfer functions by presenting predefined transfer functions for the user to choose from. However 1D transfer functions have limited classification power. Levoy [6] introduced 2D transfer functions by additionally considering gradient magnitude. This idea was taken further by the extensive work of Kniss et al. [4]. In these works 2D transfer functions based on scalar and gradient magnitude are used to identify material boundaries which are represented as arcs in the transfer function space. Sereda et al. [11] design a transfer function for easier identification of material boundaries using the LH histogram. The LH histograms represent boundaries between materials as separable clusters of points. The separability of both the arcs in 2D transfer functions and the clusters of LH histograms require materials with low internal variance. Our approach does not. In fact high and differing variances within materials can be used to discern them. This is due to our statistical approach where we consider the properties of variably sized groups of voxels instead of single voxels in isolation.

*e-mail:daniel@cmr.no
†e-mail:haidacher@cg.tuwien.ac.at
‡e-mail:Jean-paul.Balabanian@ii.uib.no
§e-mail:groeller@cg.tuwien.ac.at

Another difference is that we identify materials based on internal material properties instead of material-material interfaces as the 2D transfer function and the LH histograms do.

Other statistical approaches use ideas related to ours. The early work by Laidlaw et al. [5] considers small neighborhoods around each voxel for calculating statistical properties used for classification. Sato et al. [10] calculate a set of derived features for identifying local structures. Caban and Rheingans [1] divide the data into blocks and calculate a wide range of textural properties for each block. Similarly in the work of Lundström et al. [7], histograms of blocks are considered to improve material classification in medical data. All of these works use static neighborhoods and often calculate complex neighborhood properties. We use a dynamically changing neighborhood to better find an optimal set of voxels for calculating properties on. This approach gives good results with noncomplex neighborhood properties.

Both the works by Hadwiger et al. [3] and Correa and Ma [2] deal with size transfer functions that classify regions based on their sizes. Both methods require a preclassification of the data for defining the regions to apply the size transfer functions. For preclassification Hadwiger et al. perform region growing and calculate the volumetric size of regions whereas Correa and Ma perform diffusion by Gaussian smoothing to calculate region thicknesses. Correa and Ma also include a good overview of earlier work in scale-space analysis.

3 THE THEORY OF MOMENT CURVES

For each voxel in a volume we calculate a sequence of values based on the voxels in its vicinity. We then assign optical properties to the voxel based on this sequence of values. This results in a classification of materials or material boundaries. To calculate the sequence of values for a voxel v, we consider all its neighbor voxels that are inside a sphere centered at v with radius r (see Figure 1). For each sphere of radius r, we calculate the mean and the variance of the voxels inside. With r increasing from zero to a predefined maximum value we get a sequence of (mean μ, standard deviation σ) pairs for the voxel v. We depict this sequence of the first and second moment parameterized by r, as a curve in 3D space. As seen to the right in Figure 1 the axes of the 3D space are the mean value μ in red, the standard deviation σ in green, and the radius r in blue. Each voxel in the volume is thus represented as a curve in 3D. We call the curves moment curves since they describe the evolution of the first and second statistical moments as the neighborhood increases. By selecting certain groups of curves we end up with a classification of voxels. In the following paragraphs we describe how the shape of these curves relate to material properties.

We demonstrate the properties of the moment curves by considering a simple dataset consisting of two materials planarly touching (see Figure 2) and having infinite extent on their respective sides. The intensity values of each material are normal distributed with defined means and variances. The sequence of mean and standard deviation pairs for voxels inside a material will estimate the material's mean and standard deviation with increasing accuracy as the radius grows. We will now separately consider how the standard deviation and how the mean of a voxel evolves as the radius increases. Figure 2 shows the two materials and voxels at different positions in

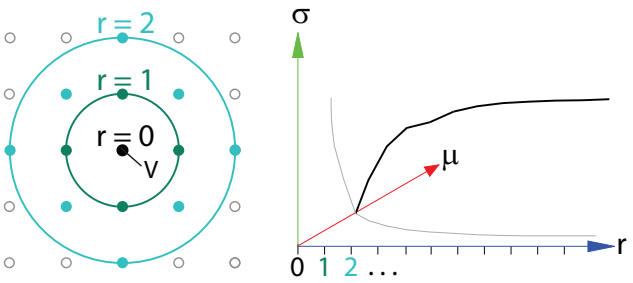

Figure 1: Left: The neighbor voxels for different radii of a voxel v. Right: A moment curve in 3D space as a function of radius r. The gray lines are the projections of the curve into the $r = 0$ plane and the $\sigma = 0$ plane.

the materials. For some voxels several circles are shown to indicate varying radii.

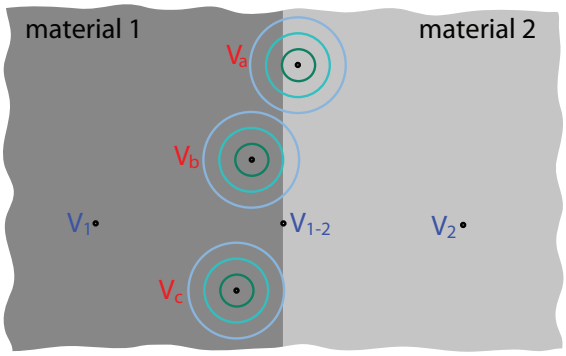

Figure 2: Two materials with different distributions touching each other.

For voxel V_1 we have plotted the evolution of the standard deviation for increasing radii. This is the curve labeled V_1 in Figure 3. As soon as the neighborhood is large enough, the curve stabilizes at the standard deviation of the material. We have colored the part of the curve after stabilization in blue and the part before in gray. A single sample has no standard deviation so all curves start at $\sigma = 0$ for $r = 0$. As the neighborhood increases, the standard deviation increases until it stabilizes. At what radius the curve stabilizes depends on the 'patchiness' of the material. The patchier it is, the higher radius it stabilizes at. A corresponding curve is plotted for voxel V_2 which is in material 2. Material 2 has lower variance than material V_1 therefore the curve stabilizes below the curve for V_1.

In Figure 4 the curve for voxel V_1 is drawn with its mean value as a function of the radius. Also here, for large enough radii, the curve stabilizes at the mean value for the material. However when the radius is 0, the mean will be the intensity of the voxel itself which can be any value from the material's normal distribution. The range of the values of material 1 is drawn as a thick vertical gray line along the μ axis. The material's distribution is drawn as well. The curve for voxel V_1 may thus start at any point along this range with a probability corresponding to the material's distribution. From the starting point, the curve will move towards the mean of the material. This is represented by a gray fan-in from the material range into the blue stabilized curve. The μ curve for voxel V_2 in material 2 is also drawn. Material 2 has a higher mean value than material 1. The distributions of the materials are overlapping. This is depicted in the overlapping vertical lines along the μ axis. It is thus possible that a curve for a voxel in material 1 intersects a curve for a voxel

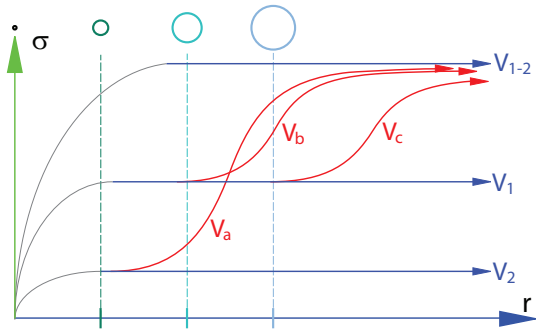

Figure 3: The standard deviation of the voxels in Figure 2 as a function of increasing radius.

in material 2 before the curves stabilize. This is represented by intersecting lines from the fan-ins of the two curves.

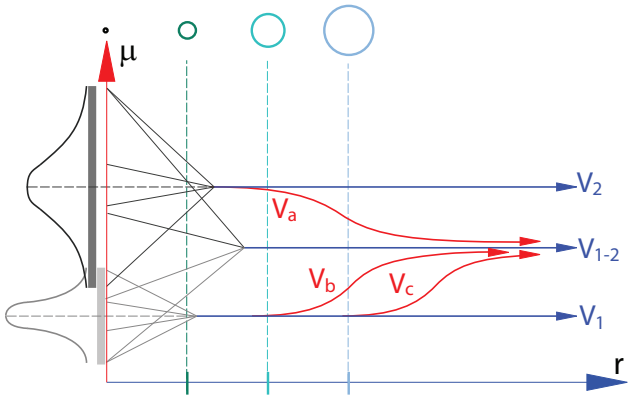

Figure 4: The mean values of the voxels in Figure 2 as a function of increasing radius.

Considering the stabilized parts of the μ and σ curves, one can separate materials with differing means or standard deviations. In contrast 1D transfer functions cannot discern materials solely based on variance and cannot binary discriminate materials with overlapping intensity distributions. Note that our curves include the 1D transfer function as a special case when assigning optical properties to voxels based on their mean values at zero radius.

Now we will consider the μ and σ curves of a voxel V_{1-2} situated exactly at the border between the two materials. As the radius increases, an equal number of samples from each of the two material distributions is included. Therefore the mean will stabilize exactly between the mean for material 1 and material 2 as shown in Figure 4. The distribution for the neighborhood of V_{1-2} will be bimodal since it consists of the sum of the distributions from each material. Therefore the fan of the V_{1-2} curve covers both the material ranges. The standard deviation for curve V_{1-2} is seen in Figure 3. This curve will stabilize at a standard deviation which is less intuitive to understand than the mean. We will give an analytical expression for this further on.

In Figure 2 we have drawn voxels V_a, V_b and V_c having different distances to the border. Their corresponding σ and μ curves are drawn in Figure 3 and Figure 4 respectively. Three circles with different radii are drawn around each of the voxels and these circles are also marked along the radius axes in the two plots. From a specific radius onward, the neighborhood will include voxels from the neighboring material. The μ and σ curves will then move away

from the stable part and will asymptotically approach the respective mean and standard deviation of V_{1-2}. This happens because the number of samples from material 1 and material 2 will equalize as the radius goes towards infinity. We will now present expressions of the mean and of the variance as a function of the ratio κ of samples between material 1 and 2. Let the two materials have mean and standard deviation (μ_1, σ_1) and (μ_2, σ_2) respectively. The derivation of these equations are given in the Appendix.

$$\mu_\kappa(\kappa) = (1-\kappa)\mu_1 + \kappa\mu_2 \tag{1}$$

$$\sigma^2_\kappa(\kappa) = -(\mu_1-\mu_2)^2\kappa^2 + \left((\mu_1-\mu_2)^2 - ({\sigma_1}^2 - {\sigma_2}^2)\right)\kappa + {\sigma_1}^2 \tag{2}$$

The mean is a linear interpolation between the μ's weighted by κ, and the variance is a second degree polynomial in κ where the coefficents are expressions of the σ's and μ's.

We will now present the relation between a border voxel at different distances to the border of the two materials and the μ and σ for the samples inside a fixed radius around that voxel. The materials can then be represented as two points in a diagram with μ on the horizontal and σ on the vertical axis. As the border voxel moves from one side of the border to the other, the ratio of voxels inside a neighborhood of fixed radius will go from material 1 ($\kappa = 0$) to material 2 ($\kappa = 1$). By using the ratio as an argument into equation 1 and equation 2, the $(\mu_\kappa, \sigma^2_\kappa)$ point for the border-voxel will move on a parabolic arc from material 1 to material 2. The square root of a second order polynomial with negative factor of the second order term, as given in equation 2, is an ellipsoid. Therefore, the $(\mu_\kappa, \sigma_\kappa)$ point will move from material 1 to material 2 along an ellipsoid as the ratio changes. Exactly at the border, the ratio κ equals 0.5 and the mean and variance is here:

$$\begin{aligned} \mu_{1-2} &= \mu_\kappa\left(\tfrac{1}{2}\right) = \tfrac{\mu_1+\mu_2}{2} \\ \sigma^2_{1-2} &= \sigma^2_\kappa\left(\tfrac{1}{2}\right) = \left(\tfrac{\mu_1-\mu_2}{2}\right)^2 + \tfrac{{\sigma_1}^2+{\sigma_2}^2}{2} \end{aligned} \tag{3}$$

When knowing the mean and variance of two bordering materials, equations (3) allow to calculate the shape of the arc that appears for voxels near the material's border. One can then automatically create a transfer function that performs a brushing which classifies this border. We showed earlier that voxels inside materials have moment curves that stabilize. We have now shown that voxels on the boundaries between two planarly touching materials also have stabilized curves. For voxels close to boundaries the moment curves switch at increasing radius from the material moment curve to the border moment curve. The radius where the transition happens indicates the distance of the voxel to the boundary. By selecting voxels that have stabilized curves which switch at a certain radius, we can identify regions with specific distances to boundaries.

4 Synthetic Data

We will now demonstrate the properties of the moment curves on a 3D test dataset consisting of 6 cylinders, each with a radius of 32 voxels, a height of 40 voxels and noise added so as to get different standard deviations. The cylinders are named cyl_σ where σ is the standard deviation in the interior.

The top left image of Figure 5 shows a slice through the cylinders. In Figure 6 we have plotted 25000 moment curves from the interior of cyl_2 and cyl_{20}. The curves are projected onto the right side and the bottom of the plot. The gray semitransparent rectangle is a cutting plane at $r = 1$. Red points are drawn on the cutting plane at positions where the moment curves intersect. One large cluster of points can be seen on this plane for cyl_{20} and one small cluster for cyl_2. The green curve shows the curve for one particular

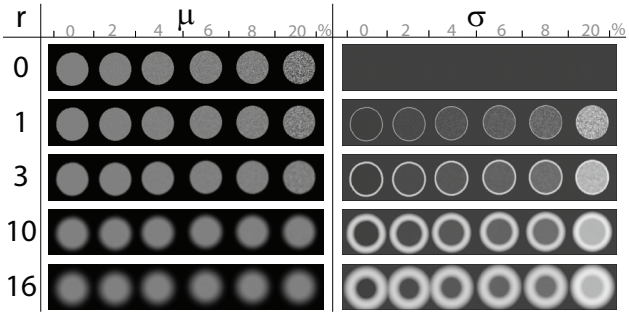

Figure 5: To the left μ and to the right σ for voxels with increasing radii are depicted. All cylinders have a $\mu = 0.5$ and an increasing $\sigma = \{0\%, 2\%, 4\%, 6\%, 8\%, 20\%\}$ from left to right.

voxel in cyl_{20}. One cluster of curves for each material is generated and each cluster converges to the mean and standard deviation of the corresponding material. The curves for the materialin cyl_2 with the low standard deviation converge faster and voxels in cyl_{20} have values from the whole intensity interval.

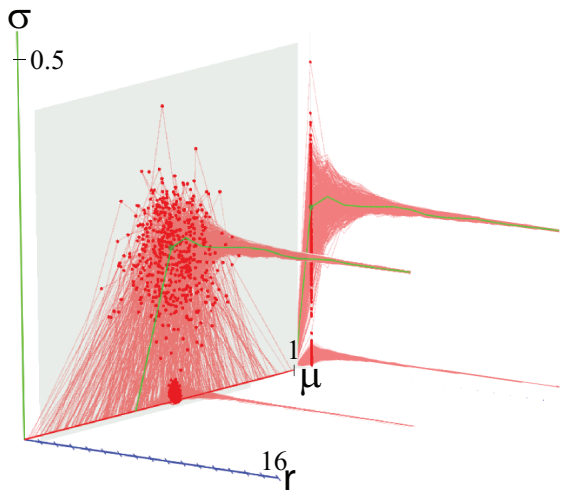

Figure 6: Moment curves, projections and $r = 1$ intersections for cyl_2 and cyl_{20} in Figure 5.

Now we will consider how moment curves of voxels close to material borders behave in practice. In Figure 7b we have drawn 31 moment curves for voxels along a line from the interior of cyl_8 to the exterior (see Figure 7c). Figure 7a and 7d show the projection of moment curves onto the r-μ and the r-σ plane. The green curve represents the voxel closest to the border. Notice how curves jump out of their stabilized path as soon as the neighborhood is large enough to cross the border. The moment curve of the border voxel is stabilizing fastest and the other curves, according to their distance to the border, are asymptotically moving towards it. In Figure 7b) we show a cutting plane at $r = 16$ with curve intersections marked as red dots. Since we know the σ's and μ's of each material, we can get an expression for μ and σ for voxels moving from material 1 to material 2. By substituting $\mu_1 = 0.5$, $\sigma_1 = 0.08$ for cyl_8 and $\mu_2 = 0$, $\sigma_2 = 0$ for the exterior into equation 2 we get $\mu_\kappa = 0.5\kappa$ and $\sigma^2_\kappa = -0.25\,\kappa^2 + 0.2564\,\kappa$. The analytic $(\mu_\kappa, \sigma_\kappa)$ path overlaid on the intersection points at $r = 16$ is depicted in Figure 8. The tight overlap demonstrates the expected correspondance between

the analytic formulas and the discrete samples of the test dataset.

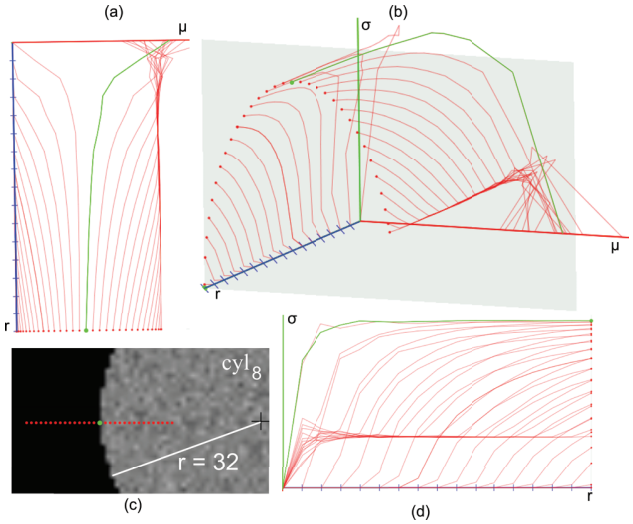

Figure 7: Moment curves for voxels along a line from the inside to the outside of cyl_8. The green curves and point represent a voxel on the boundary of the material.

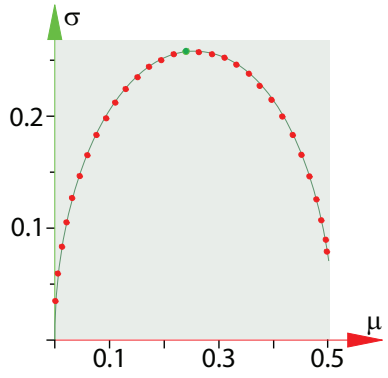

Figure 8: The curve intersections at $r = 16$ as red dots and an analytically plotted (μ_κ, σ_κ) path overlaid to show correspondence.

We have now discussed the different types of moment curves that occur in our dataset. In Figure 9a we plot moment curves representing voxels from the entire volume. One can clearly see clusters for each cylinder and arcs from each material to their surrounding background values. We apply a transfer function on the moment curves to classify the cylinders and their borders. First we try to classify cyl_8 by defining a blue rectangle on a cutting plane at $r = 16$. The rectangle is located around the cluster for cyl_8. Our transfer function is defined such that each voxel having a moment curve that intersects the rectangle at $r = 16$ will be classified and colored. The result can be seen at the bottom of Figure 9a.

The interior voxels of cyl_8 which have a distance of more than r=16 voxels to the border are now classified in blue. Voxels with distance less than 16 to the border are not classified because their curves will have 'jumped' away from the stabilized path at earlier radii. The transfer function also classifies circles inside each of the cylinders of lower variance. The over-classified voxels are due to moment curves which have 'jumped' away from their stabilized part and intersect the rectangle as they move toward the top of the arcs between material and exterior. We can avoid classifying these regions by only classifying curves that are stabilized as they intersect the rectangle. The degree of stabilization of a moment curve can be expressed by its partial derivative along r. We define the derivative curve as having it's μ' value for a certain radius r equal to the difference between the μ at $r-1$ and r. Similarly we find the curve's σ' value for a radius r to be the difference between the σ at r-1 and r. These derivative curves can be visualized in the same manner as the moment curves. In Figure 9b we are visualizing all the derivative curves. This yields an interesting pattern where the curves seem to move along a cylinder which has a decreasing thickness along the r axis.

We can select moment curves with an arbitrary degree of stabilization by brushing on the derivative curves. In Figure 9c the (μ',σ') pairs for each curve intersecting $r = 16$ can be seen. The blue rectangle defines the curves we want to brush to select stabilized moment curves.

In Figure 9d the moment curves corresponding to the points brushed in the derivative space are shown in blue. Only stabilized curves are in blue and these stabilized curves are found in the interior of each material, at the border between each material and background, and in the background itself. We now assign blue to any voxel having a stabilized curve according to the brush in the derivative space in Figure 9d. The transfer function is seen as a rectangle covering all moment curves of the data. By selecting all stabilized curves we get the interiors and the borders of the materials as well as the background (see bottom of Figure 9d). Close to the borders safety margins of not-selected voxels appear. The width of these margins depends on the chosen radius r.

Now we continue to refine the transfer function with different colors for each easily separable cluster of stabilized curves. The transfer function is defined as a set of coloured polygons in the plane at $r = 16$ (top of Figure 9e). The corresponding classification is seen at the bottom of Figure 9e.

After brushing the curves we are able to select for example only cyl_8. In the top row of Figure 10 we do a volume rendering of the test dataset using a semitransparent 1D black-to-white transfer function and with cyl_8 classified as blue. Due to the safety margin the classified cylinder is too small. To compensate for the shrinking, we dilate the volume r voxels where r is the radius of the plane the transfer function was defined on. In the bottom row of Figure 10 an image series of the dilation steps for the cylinder can be seen. At the last step it almost covers its original space again. However the edges are rounded, so a perfect reconstruction has not been made. By defining transfer functions at cutting planes with lower radii one gets less rounding of edges. This has the disadvantage of brushing less stabilized curves since the maximal radii considered are smaller.

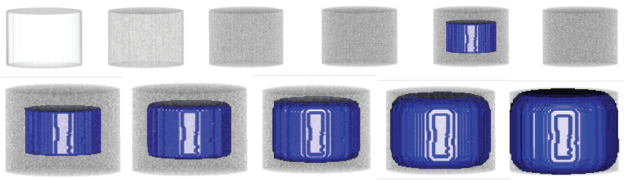

Figure 10: Top row shows volume rendering with cylinder cyl_8 classified. The bottom row shows the dilation steps for the cylinder.

We propose the following workflow for the transfer function design. First select a radius where most materials have stabilized curves ($r = 16$ in our example). Visualize the intersection of the curves with a cutting plane at the selected radius. This reduces the 3D space to a 2D plane where it is easier to define a transfer function. Depending on the data, the 2D scatterplot can be cluttered and contain many arcs. Stabilized paths of the scatterplot can

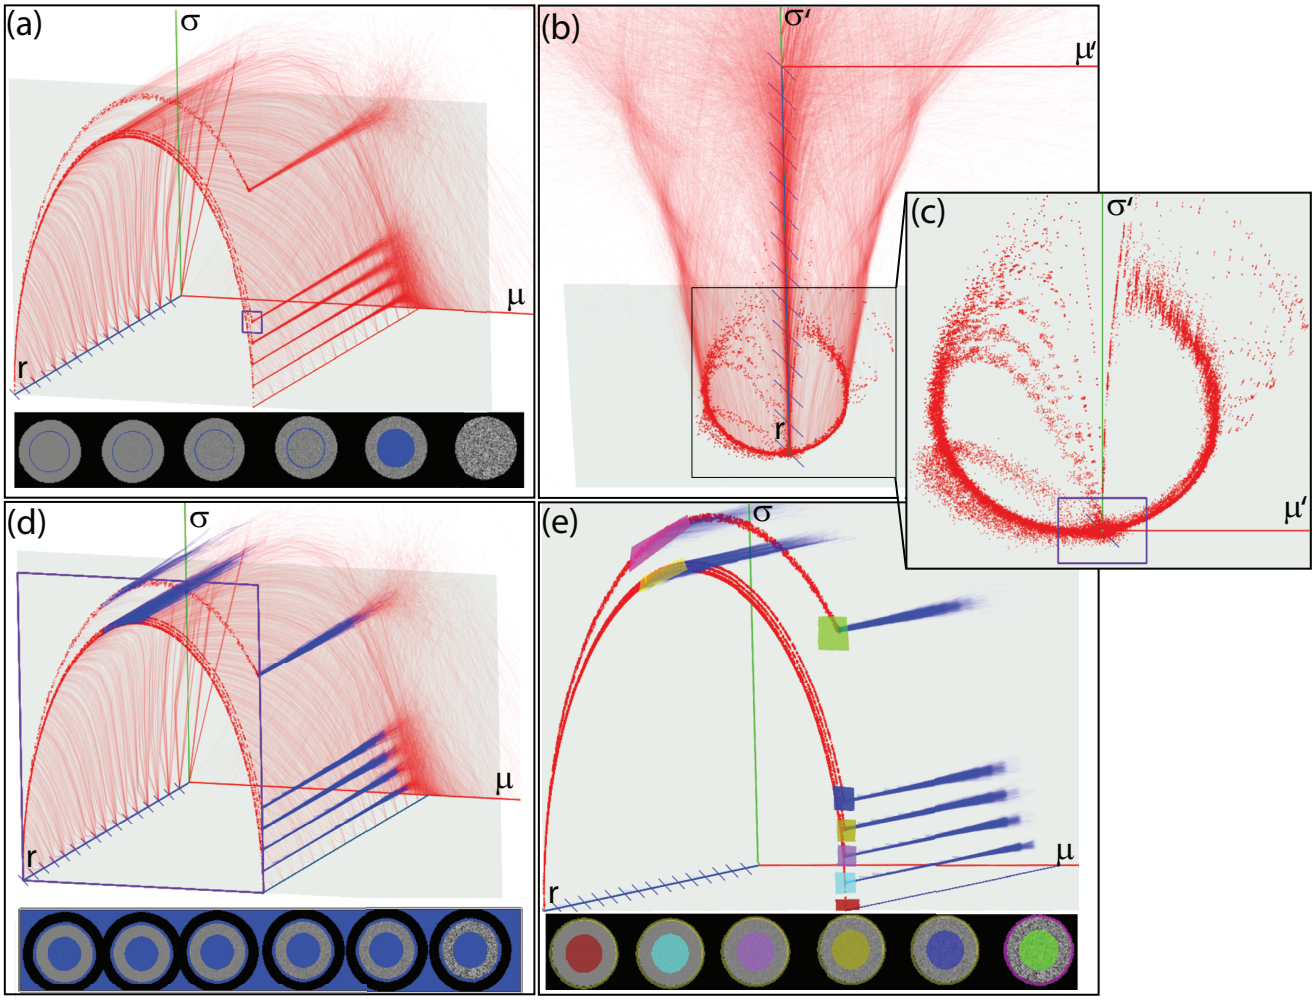

Figure 9: (a) The moment curve space of the cylinder test data with a blue brush rectangle and the corresponding classification below. (b) Derivative curves from the moment curves. (c) The derivative scatterplot of the moment curves intersecting the plane $r = 16$. (d) The curve space of the test dataset with stabilized curves shown in blue. A derivative brush as defined in c) selects all stabilized curves. The result is a classification of all material interiors, material borders, and background. (e) A color transfer function on stabilized moment curves identifies materials of different variances and their borders.

be identified by brushing the corresponding derivative scatterplot (Figure 9c). Now one can assign colors to the different clusters of stabilized curves. Figure 11a shows the projected moment curves in red and the stabilized clusters in blue. Figure 11b shows only stabilized curves yielding a clear view with easily identifiable regions. The zoom-in corresponds to the material boundaries of $cyl_{0,2,4,6,8}$.

5 Implementation

The classification procedure is performed in four steps. Firstly the user selects an area on a slice positioned through the volume. The moment curves for the voxels in that area are calculated on the CPU and visualized. Curves for areas on other slices can be added. Secondly the user performs a brushing on the moment curves. A classification is then performed based on the brush on all samples on the slice. By exploiting the parallelity of the GPU we are able to do this in real time. The frame rate is sensitive to the maximum calculated radius. The number of voxels being considered for each fragment of the slice increases cubically as a function of the radius. Using precomputed means and variances could considerably reduce the classification time but at the expense of texture memory. Thirdly, as soon as the user finds the slice classification satisfactory, a 3D classification volume is built by iteratively classifying each slice in the volume. On an Nvidia Geforce 8800 GTS graphics card this can take from a few seconds on small datasets up to a minute on the CT dataset of size $512 \times 512 \times 1112$, as used in the results Chapter. Finally, a distance transform up to a distance of r voxels on the classified volume is performed. Afterwards we render in real-time the classified regions and dilate them up to r voxels by changing the isovalue of the region borders. The distance field from the dilation step is additionally used in calculating normals for Phong shading the classified regions. Our method is different from seeding and growing approaches for data segmentation since all voxels are classified independently from each other. Also the method is global and therefore does not require a specific seeding.

6 Results

We have applied our transfer function successfully on a CT dataset of a human with a resolution of 512x512x1112 and on an MR dataset of a sheep heart with a resolution of 352x352x256. We follow the same approach as suggested in Chapter 4 for the cylin-

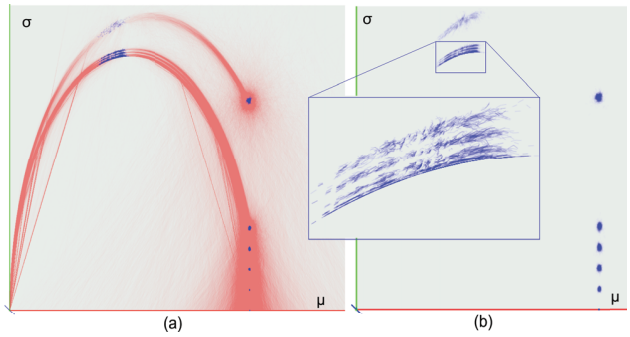

Figure 11: (a A 2D projection of the 3D moment curve space with identified stable regions in blue. (b Only stabilized curves are shown with a zoom-in on the boundaries.

der test dataset. The steps for analyzing the CT torso can be seen in Figure 12. The moment curves of 30000 evenly spaced samples taken on the leftmost slice in 12g are shown in 12a. To find stabilized curves we calculate the derivative curves 12b and brush (small blue rectangle) in the r=16 cutting plane in 12c. After brushing the stabilized moment curves are visualized in blue in 12d. For easy brushing we use the 2D projection of all curves in 12e. In 12f we define coloured brushes on easily identifiable regions. The resulting classification for three different slices is shown in 12g. Muscles are colored green, fat is colored blue, the liver is red, and the spleen is bright cyan. Also some voxels in the lungs are classified separately as dark cyan. Finally the transition between fat and muscles is pink. We dilate the spleen and liver and volume render them in Figure 13. For the background material a semitransparent 1D black-to-white transfer function is used.

The second dataset we use is an MRI of a sheep heart. This dataset is challenging to create meaningful classifications and has therefore been used in transfer function papers as a benchmark. In the paper by Pfister et al. [9] several 1D transfer function variants as well as the 2D transfer function by Kniss et al. [4] (Figure 14a) were tested on this dataset. Also the LH histogram has been applied on this dataset by Sereda et al. [11](Figure 14b). From the comparison in Pfister et al. [9] and even by quickly experimenting with a 1D transfer function it follows that a 1D transfer function only manages to separate the background from the sheep heart but fails to clearly separate interior materials in the heart. The LH histogram [11] attempts to separate a higher valued area in the middle of the heart from the rest. However by looking at the LH transfer function one can see that the yellow and red part extend over almost the entire F_H domain and therefore basically define a 1D transfer function which will fail to separate the middle part from the rest. Also the 2D transfer function by Kniss et al. [4] used in Pfister et al. [9] has a continuous color gradient and thus does not clearly separate the two regions. Our approach also failed to classify the middle region. This is partly due to the small size of the dataset with moment curves that cannot achieve a big enough radius to stabilize before they grow out of the region. It is also due to a lack of homogeneity of the middle region. The darker colored middle region has bright thin structures going through it. This disables curve stabilization. To address these issues we apply a two level approach. First we define a simple 1D transfer function that classifies the middle region but also over-classifies some areas outside it. Then we calculate the moment curves by only considering voxels with values in the interval defined by the 1D transfer function. We have defined the 1D transfer function to exclude the thin structures that were inside the region we try to classify. Then also the neighboring materials are excluded. Now the moment curves have a better chance to stabilize. With this technique we were able to classify the middle region as can be seen in Figure 14c. The slices in Figure 14d-h show the classified areas in green and the areas within the 1D transfer function interval in red. We believe we have comparable or better results with our approach than earlier published results as can be seen in the slices and the volume renderings.

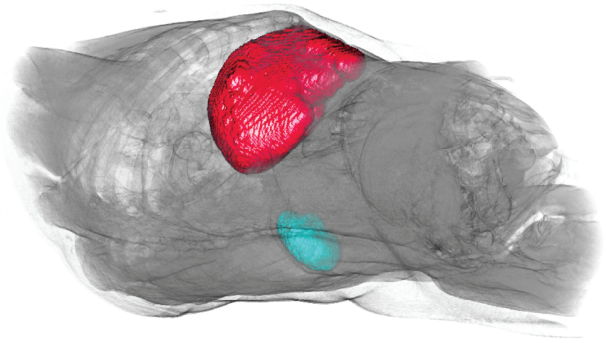

Figure 13: Two classified and dilated regions in a human CT dataset.

7 Conclusion and Future Work

We have introduced moment curves and showed how they behave in different material arrangements. We have described a workflow for applying a specific transfer function to classify materials and material boundaries. The necessary transfer function editing is performed in 2D on easily identifiable clusters. Finally we have applied our approach on a CT dataset of a human torso and classified separate tissue types. We have also successfully applied a slightly modified two-level version of the transfer function on MR data and compared it with existing transfer functions. We believe moment curves has potential on other types of data as well. By combining our method with an automatic cluster-identification algorithm, one can imagine a fully automatic transfer function specification. The stabilized moment curves for a volume would be calculated. The cluster-identification algorithm would segment each cluster. In the successive rendering users could assign optical properties or tags to each segmented object.

Acknowledgements

The authors wish to thank Dr. Ruben Patel, Professor Helwig Hauser and Professor Ivan Viola for fruitful discussions and also Dr. Christopher Giertsen for supporting this work.

References

[1] J. Caban and P. Rheingans. Texture-based transfer functions for direct volume rendering. *IEEE TVCG*, 14(6):1364–1371, 2008.

[2] C. Correa and K.-L. Ma. Size-based transfer functions: A new volume exploration technique. *IEEE TVCG*, 14(6):1380–1387, 2008.

[3] M. Hadwiger, F. Laura, C. Rezk-Salama, T. Hollt, G. Geier, and T. Pabel. Interactive volume exploration for feature detection and quantification in industrial CT data. *IEEE TVCG*, 14(6):1507–1514, 2008.

[4] J. Kniss, G. Kindlmann, and C. Hansen. Multidimensional transfer functions for interactive volume rendering. *IEEE TVCG*, 8(3):270–285, 2002.

[5] D. H. Laidlaw, K. W. Fleischer, and A. H. Barr. Partial-volume bayesian classification of material mixtures in MR volume data using voxel histograms. Technical report, California Institute of Technology, Pasadena, CA, USA, 1997.

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

[7] C. Lundström, P. Ljung, and A. Ynnerman. Local histograms for design of transfer functions in direct volume rendering. *IEEE TVCG*, 12(6):1570–1579, 2006.

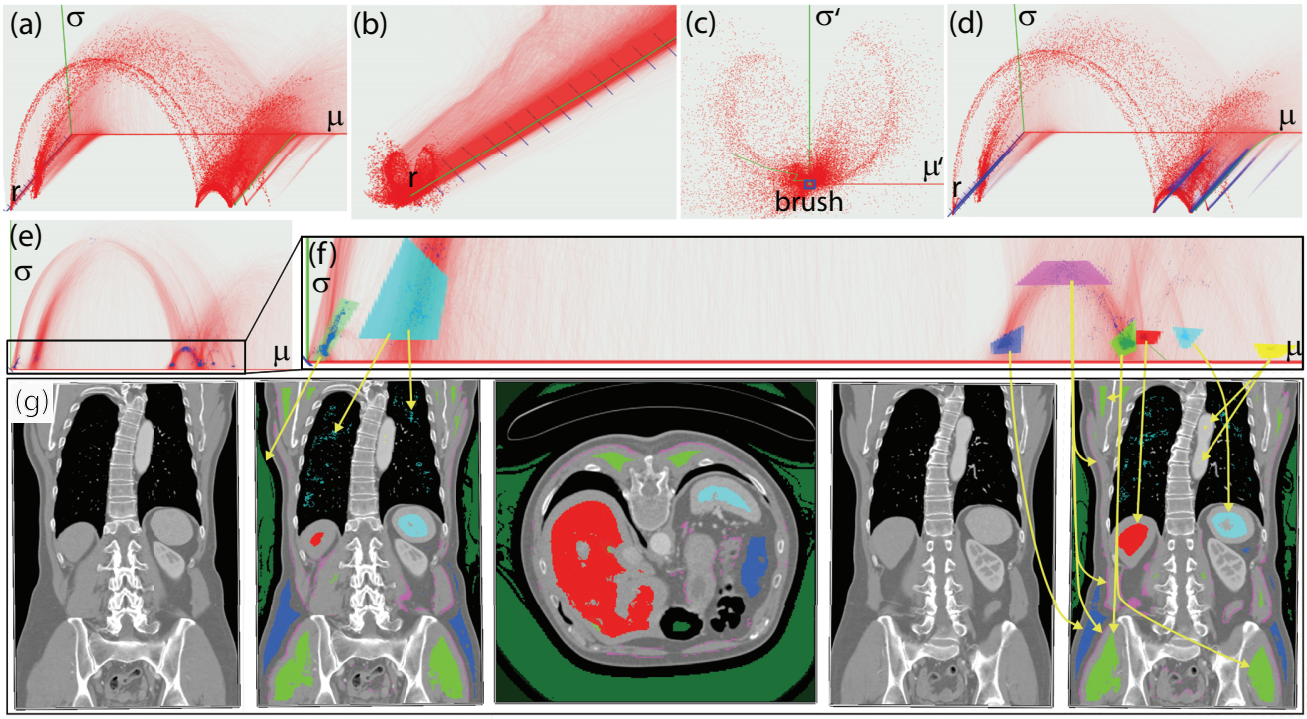

Figure 12: Moment curve steps a)-f) for classifying regions of a human CT dataset. Results show in e) two coronal slices before and after classification and one classified axial slice.

[8] J. Marks, B. Andalman, P. A. Beardsley, W. Freeman, S. Gibson, J. Hodgins, T. Kang, B. Mirtich, H. Pfister, W. Ruml, K. Ryall, J. Seims, and S. Shieber. Design galleries: a general approach to setting parameters for computer graphics and animation. In *SIGGRAPH '97*, pages 389–400, 1997.

[9] H. Pfister, C. Bajaj, W. Schroeder, and G. Kindlmann. The transfer function bake-off. *VIS '00: Proceedings of the 11th IEEE Visualization 2000*, pages 523–526, 2000.

[10] Y. Sato, C.-F. Westin, A. Bhalerao, S. Nakajima, N. Shiraga, S. Tamura, and R. Kikinis. Tissue classification based on 3D local intensity structures for volume rendering. *IEEE TVCG*, 6(2):160–180, 2000.

[11] P. Sereda, A. Bartroli, I. Serlie, and F. Gerritsen. Visualization of boundaries in volumetric data sets using LH histograms. *IEEE TVCG*, 12(2):208–218, 2006.

APPENDIX

In this section equations (1) and (2) of Chapter 3 are derived.

The equations for calculating the mean and variation for a discrete collection of samples $\{x_1, x_2, ..x_n\}$, are given in (4) and (5):

$$\mu = \frac{1}{n}\sum_{i=1}^{n} x_i \tag{4}$$

$$\sigma^2 = \frac{1}{n}\sum_{i=1}^{n}(x_i-\mu)^2 = \frac{1}{n}\sum_{i=1}^{n} x_i^2 - \mu^2 \tag{5}$$

Assume a material 1 has n_1 samples indexed $\dot{x}_i$ where the mean of the samples are μ_1 and the variance is σ_1^2. Similarly, assume a material 2 with n_2 samples $\ddot{x}_i$ having mean μ_2 and variance σ_2^2. We want to find the mean μ_3 and variance σ_3^2 for a material 3 consisting of the intensity values of material 1 and 2: $\{\dddot{x}_1, \dddot{x}_2, ..., \dddot{x}_{n_3}\} = \{\dot{x}_1, \dot{x}_2, ..., \dot{x}_{n_1}, \ddot{x}_1, \ddot{x}_2, ..., \ddot{x}_{n_2}\}$. We have:

$$n_3 = n_1 + n_2 \tag{6}$$

Since addition is associative we have:

$$\sum_{i=1}^{n_3} \dddot{x}_i = \sum_{i=1}^{n_1} \dot{x}_i + \sum_{i=1}^{n_2} \ddot{x}_i \tag{7}$$

$$\sum_{i=1}^{n_3} \dddot{x}_i^2 = \sum_{i=1}^{n_1} \dot{x}_i^2 + \sum_{i=1}^{n_2} \ddot{x}_i^2 \tag{8}$$

We now want to express μ_3 and σ_3 as a function of μ_1, μ_2, σ_1, σ_2, n_1 and n_2. By rearranging (4) for material 1 and 2 we get:

$$\sum_{i=1}^{n_1} \dot{x}_i = n_1\mu_1 \quad \sum_{i=1}^{n_2} \ddot{x}_i = n_2\mu_2 \tag{9}$$

Similarly by rearranging (5) we get:

$$\sum_{i=1}^{n_1} \dot{x}_i^2 = n_1(\sigma_1^2 + \mu_1^2) \quad \sum_{i=1}^{n_2} \ddot{x}_i^2 = n_2(\sigma_2^2 + \mu_2^2) \tag{10}$$

For μ_3, substituting (6) and (7) into (4), we get the expression in the middle of (11). By applying (9) on the middle expression we get the right side expression of (11):

$$\mu_3 = \frac{1}{n_3}\sum_{i=1}^{n_3} \dddot{x}_i = \frac{1}{n_1+n_2}\left(\sum_{i=1}^{n_1} \dot{x}_i + \sum_{i=1}^{n_2} \ddot{x}_i\right) = \frac{1}{n_1+n_2}(n_1\mu_1 + n_2\mu_2) \tag{11}$$

For the variance σ_3 we substitute (10) into (8) and apply the result on (5) to get the right hand of the first line in (12). Then by further substituting (6) and (11) we get the second line in (12):

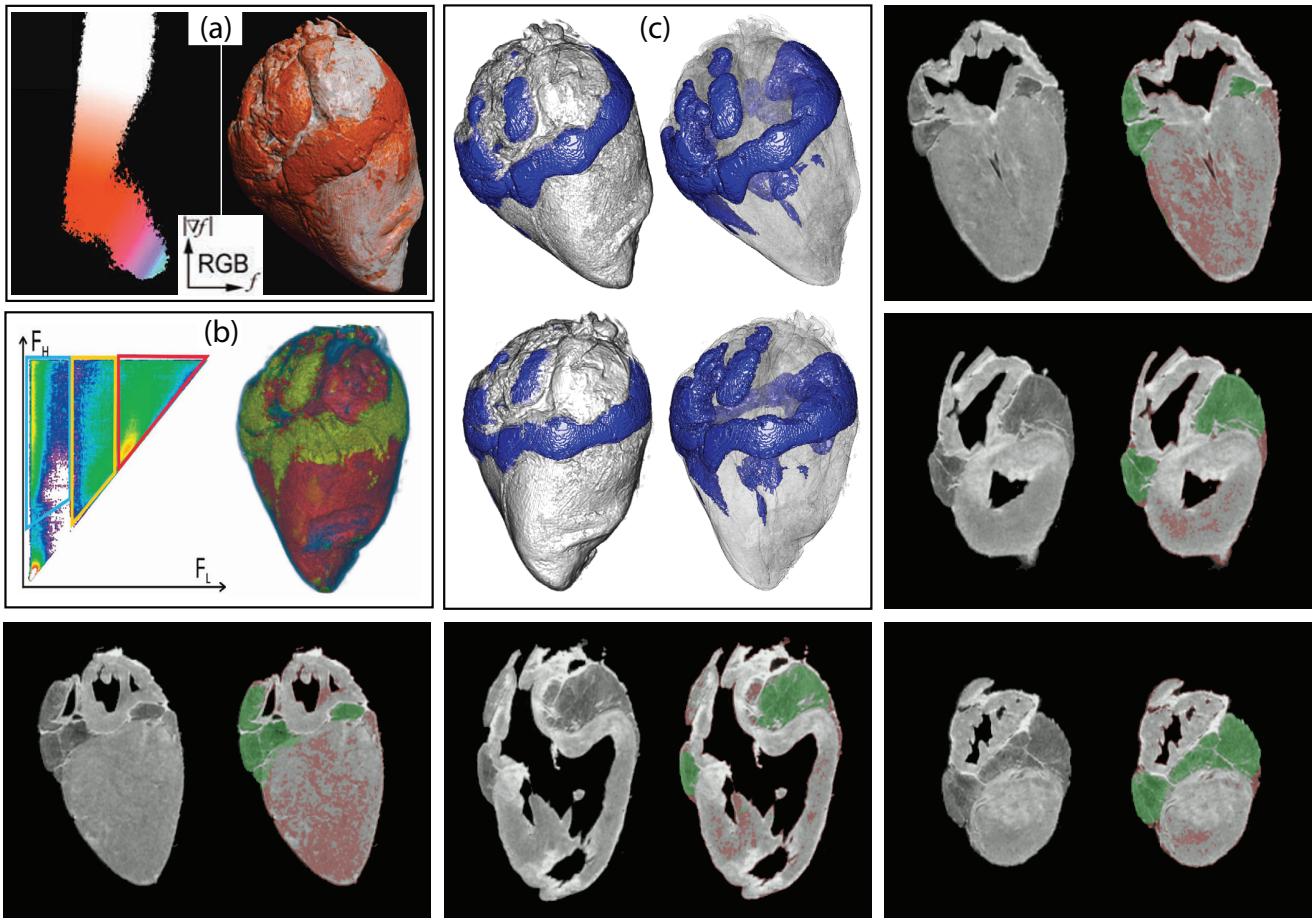

Figure 14: Comparison of (a) 2D transfer function, (b) LH histogram and (c) moment curve classification. Renderings to the right in (c) are semitransparent. Unclassified and classified moment curve slice pairs are also shown.

$$\sigma_3^2 = \tfrac{1}{n_3}\textstyle\sum_{i=1}^{n_3} \ddot{x}_i^2 - \mu_3^2 = \tfrac{1}{n_3}(n_1(\sigma_1^2+\mu_1^2)+n_2(\sigma_2^2+\mu_2^2)) - \mu_3^2$$
$$= \tfrac{1}{n_1+n_2}(n_1(\sigma_1^2+\mu_1^2)+n_2(\sigma_2^2+\mu_2^2)) - (\tfrac{1}{n_1+n_2}(n_1\mu_1+n_2\mu_2))^2 \quad (12)$$

We do not want to specify both n_1 and n_2 but rather the ratio κ between them. $\kappa = 0$ means only samples from material 1, $\kappa = 0.5$ means an equal number of samples from both materials and $\kappa = 1$ means only samples from material 2 are taken. We therefore express n_1 and n_2 as a function of κ and n_3:

$$n_1 = n_3(1-\kappa), \quad n_2 = n_3\kappa \quad (13)$$

A material 3 with ratio κ between samples from material 1 and samples from material 2 will thus have the mean $\mu_3(\kappa)$ from inserting (13) into the right hand expression of (11):

$$\mu_3(\kappa) = \frac{1}{n_3}(n_3(1-\kappa)\mu_1 + n_3(1-\kappa)\mu_2) \quad = \quad \boxed{(1-\kappa)\mu_1 + \kappa\mu_2} \quad (14)$$

This concludes the derivation of equation (1) in Chapter 3.

Similarly, a material 3 with ratio κ between samples from material 1 and samples from material 2 will have the variance σ_3^2 from inserting (13) into (12) and using (14) for the mean:

$$\sigma_3^2(\kappa) = \frac{1}{n_3}\left(n_3(1-\kappa)(\sigma_1^2+\mu_1^2) + n_3\kappa(\sigma_2^2+\mu_2^2)\right) - ((1-\kappa)\mu_1+\kappa\mu_2)^2 \quad (15)$$

We expand the parentheses in (15) and rearrange the terms so that they form a second order polynomial in κ:

$$\begin{gathered}(\sigma_1{}^2+\mu_1^2-\kappa\sigma_1{}^2-\kappa\mu_1{}^2)+(\kappa\sigma_2{}^2+\kappa\mu_2{}^2)- \\ (\kappa^2\mu_1^2+2\kappa\mu_1\mu_2-2\kappa^2\mu_1\mu_2+\mu_1^2-2\kappa\mu_1^2+\kappa^2\mu_2^2) = \\ \\ \left(-\mu_1{}^2+2\mu_1\mu_2-\mu_2{}^2\right)\kappa^2 \\ +\left(-\sigma_1{}^2+\mu_1{}^2+\sigma_2{}^2+\mu_2{}^2-2\mu_1\mu_2\right)\kappa \\ +\sigma_1{}^2 = \\ \\ -\left(\mu_1{}^2-2\mu_1\mu_2+\mu_2{}^2\right)\kappa^2 \\ +\left(\left(\mu_1{}^2-2\mu_1\mu_2+\mu_2{}^2\right)-\left(\sigma_1{}^2-\sigma_2{}^2\right)\right)\kappa \\ +\sigma_1{}^2 = \end{gathered} \quad (16)$$

$$\boxed{-(\mu_1-\mu_2)^2\kappa^2 + \left((\mu_1-\mu_2)^2 - (\sigma_1^2-\sigma_2^2)\right)\kappa + \sigma_1{}^2}$$

This concludes the derivation of equation (2) in Chapter 3.