Investigating measures for transfer function generation for visualization of MET biomedical data

L. Svensson, I. Nyström, S. Svensson and I.-M. Sintorn
Centre for Image Analysis, Uppsala University and Swedish University of Agricultural Sciences
Box 337
SE-751 05 Uppsala, Sweden
E-mail: {lennart,ingela,stina,ida.sintorn}@cb.uu.se

ABSTRACT

In this paper, the question of automatically setting transfer functions for volume images is further explored. More specifically, the focus is automatic visualization of Molecular Electron Tomography (MET) volume images using one-dimensional transfer functions. We investigate how well a few general measures based on density, gradient, curvature and connected component information are suited for generating these transfer functions. To assess this, an expert has set suitable transfer function levels manually and we have studied how these levels relate to different characteristics of the selected measures for 29 data sets. We have found that the measures can be used to automatically generate a transfer function used to visualize MET data, to give the user an approximate view of the components in the image.

Keywords

Volume visualization, direct volume rendering, transfer functions, automatic visualization, molecular electron tomography

1 INTRODUCTION

Automatic visualization provides the means for screening large amounts of data in a short time by aiding the user in setting visualization parameters. Here, the goal is to investigate measures for automatically creating one-dimensional transfer functions that give good first renderings of Molecular Electron Tomography (MET) data. These should highlight the most important information, i.e., the molecular surface of proteins, and still show other variations in the imaged sample. The visualization should be a starting point for interactive adjustments. Primarily, the focus is to identify measures which generate an appropriate opacity function.

MET allows for studying the structure and flexibility of molecules and macromolecules in solution (in vitro) as well as in tissue samples (in situ). The imaging technique reveals material density with a resolution as low as a few nanometers. For determining how the proteins function in their natural environment, tissue samples are analyzed directly using MET.

For determining molecule flexibility and dynamics, molecules in solution are analysed, see Klaile [5] for a recent example. In order to investigate, explore, and analyse the complex data, adequate visualization of the data is required.

For general automatic transfer function generation, the "transfer function bake-off" [7] presents four approaches: (1) trial-and-error, (2) data-centric without model, (3) data-centric with model, and (4) image-based. Multi-dimensional transfer functions based on curvature have been introduced by Kindlmann [4], transfer functions specified as the sum of Gaussians were presented by Kniss [6], schemes based on topology differentiation have been suggested [11, 12], Rezk Salama presented an approach [8] which focuses on parameters relating to the user's domain knowledge. The mentioned methods move transfer function generation close to identification and segmentation problems.

To our knowledge, no transfer function generator tailored for MET volumes has so far been suggested. The volumes are usually rendered with direct volume rendering with a 1-D transfer function that is manually set. Often, the pre-integration step, that was presented by Engel [2], is left out, leading to notable visualization artefacts. The transfer function generation problem is highly relevant for this type of data since the volumes are difficult to interpret, see Figure 1, due to low contrast, small objects, missing data, etc. Another problem with this type of data is that many factors af-

fect the density values in the volumes. Among the influencing factors are the energy of the electrons hitting the sample, the sample thickness and composition, the detector, and the MET reconstruction algorithm.

We address this issue by following the second approach in the bake-off paper [7], the "data centric without model" approach. The idea is to make approximate distinctions between objects using measures in the density range of the volume image. The aim is to establish relations for generating a transfer function in this domain automatically. We have chosen a one-dimensional transfer function because of the intrinsic property that densities are remapped to opacities in a consistent way. Multi-dimensional transfer functions can be useful to differentiate between regions, but there is a less clear connection to the underlying density. If a model-based approach was employed, where individual components would be identified using various means of thorough image analysis, it would have enabled more elaborate fine-tuning of the visualization. This would however be at the expense of having a more complex solution and time-consuming algorithm, as well as possibly lower generality.

Using measures in the density range is similar to the approach presented by Bajaj [1], but the calculated functions are different and are also suggested to be used in a different way. We focus on automatic extraction of isovalues, whereas Bajaj suggested his measures to be used for interactive isovalue selection and volume data exploration. We have studied four different measures based on density, gradient, curvature and connected component functions. The interesting point is how features of these measures relate to manually chosen levels by an expert. For the density histogram, it has been investigated what percentiles the manually set levels correspond to, whereas for the gradient, curvature and connected component measures, different features of the functions have been correlated to the manually set level. We have used 29 volumes in the tests.

2 IMAGE DATA

In MET, an electron microscope is used to capture 2D micrograph images from different angles of a sample. This results in a so called tilt series. The sample is a very thin frozen or chemically embedded slice. A back projection technique is applied to reconstruct a 3D image of the sample, which is then refined in an optimization procedure resulting in a *MET volume* [10]. Its scalar values correspond to the density of the sample. The complete process from sample preparation to a final volume is a long and tedious process, so these volumes are not available in large quantities.

MET volumes are difficult to interpret for a number of reasons: the resolution is relatively low — each protein is represented by a small number of voxels; the contrast is low as electron irradiation destroys the sample, which means that the total dose used to acquire the micrographs must be kept low; the MET volume suffers from missing data artefacts as the electron microscope limits the angular range to 120° – 140°; and the density levels in the MET volume are relative and not absolute. For an untrained eye, the volumes often seem to only contain a large number of blobs of varying size and shape.

Data set	No of invest. volumes	No of molecule instances per vol.
IgG	3	$\sim$2
RNAP II	3	$\sim$10
CEACAM1	6	$\sim$90
TMV	17	$\sim$2

Table 1: Investigated MET volumes

Another challenge when visualizing MET volumes compared to, e.g., MRI volumes, is that the molecules cannot be studied individually using visible light. This means there is no ground truth to refer to regarding how they should be visualized.

We investigate and evaluate the chosen measures using 29 MET volumes from four different studies of proteins in solution. All MET volumes were reconstructed using the constrained maximum entropy tomography method [10], giving a few nanometer resolution. The proteins are described in brief below. Since this kind of data is difficult and labour intense to generate, 29 volumes should be considered as a relatively large set of volumes. See Table 1 for a list of the number of respective protein images and molecules.

The RNA Polymerase II (RNAP II) is a large, fairly round macromolecule responsible for mRNA synthesis in eukaryotic cells. In the investigated MET volume, a RNAP II macromolecule has a $\sim$ 21 voxel diameter.

The Immunoglobulin G (IgG) antibody is a smaller macromolecule. Antibodies are crucial parts of our immunological defence system, e.g., IgG binds to foreign agents such as virus particles and targets them for destruction. An IgG antibody has three roundish parts of equal size connected at one center point. Two of the arms are fragment antigen binding arms and one is a fragment crystallisable stem. In the investigated MET volume, the smallest round subpart has a $\sim$ 10 voxel diameter. See Sandin [9] for details.

The carcinoembryonic antigen related cell adhesion molecule 1 (CEACAM1) is even smaller than IgG. It is a transmembrane receptor involved in binding with other cells. In the investigated MET volumes, CEACAM1 occur as monomers, one molecular unit, with a volume of $\sim$ 580 voxels, or dimer, two molec-

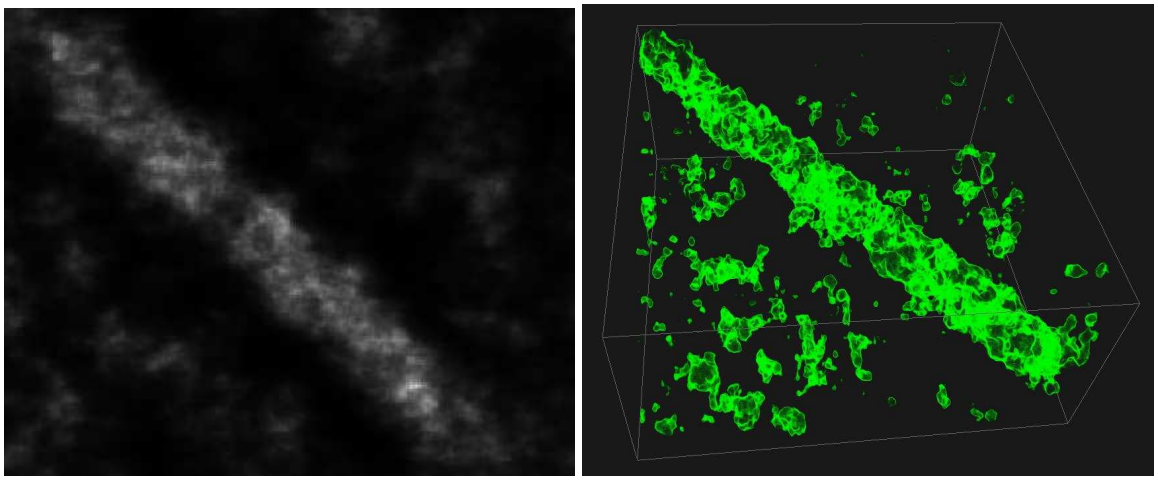

Figure 1: The Tobacco Mosaic Virus (TMV) reconstructed using Molecular Electron Tomography (MET). Left: a slice of a MET volume is shown. Right: a volume visualization with a manually set transfer function.

ular units linked together, with a volume of ~ 1160 voxels. See Klaile [5] for details.

The Tobacco Mosaic Virus (TMV) is a larger, tubular structure. In the investigated MET volume, the TMV has ~ 32 voxel diameter. It infects, e.g., tobacco plants.

The MET volumes in the study are approximately 256^3 voxels. The protein concentrations are such that we expect ten RNAP II molecules, one or two IgG antibodies, and 80–100 CEACAM1 molecules in each MET volume. For TMV, one to three structures are present in each volume.

3 INVESTIGATION

A MET volume can be described as $f : \mathbb{N}^3 \rightarrow \mathbb{R}$, not considering the limitations of digital number representation, while the used transfer function corresponds to $g : \mathbb{R} \rightarrow \mathbb{R}^4$. The four output components of the transfer function are the three color channels, RGB, and the opacity.

We suggest to render a MET volume with a transfer function built using two primitives, a Gaussian and a piece-wise linear function going from transparent to opaque, see Figure 2 (top). The idea behind the Gaussian is that it should correspond to the most central surface level for the molecules of interest. This we denote the *primary level*. With a high concentrated peak, the Gaussian will have a similar effect as an iso-surface, but both reveal the shape of objects and give an idea about the blurriness around it, i.e., if the surface is part of a sharp change or a smooth transition. The piece-wise linear function has two roles, partly to reveal more of the MET volume and partly to highlight outliers. To visualize this, a linear ramp with low opacity slope is suggested, starting at the *secondary level*, to not disturb the surface visualization, while a linear ramp with high slope in a different color is recommended for the outlier density range, starting at the *outlier level*.

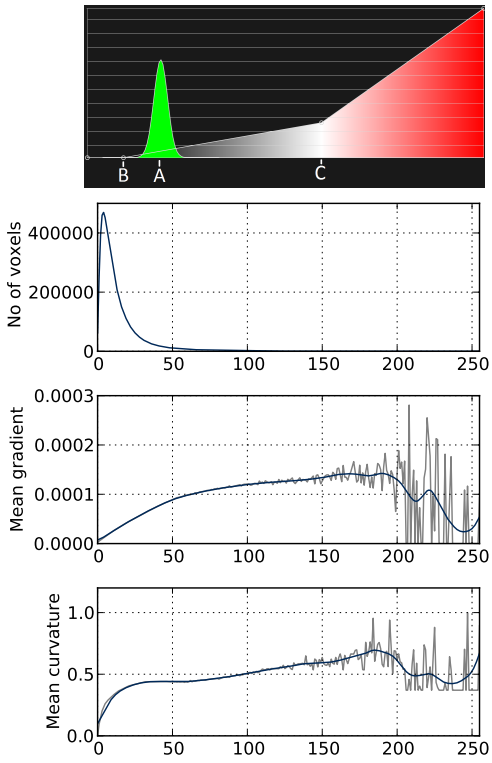

Figure 2: Example of transfer function (top). Plots of the histogram and the gradient and curvature functions of a MET volume containing a Tobacco Mosaic Virus (bottom).

To set the positions for these primitives, the following estimates need to be extracted from the chosen measures:

- Primary level (center of the Gaussian, 'A' in Figure 2)
- Primary level standard deviation (width of the Gaussian)
- Secondary level start (start of low slope ramp, 'B' in Figure 2)
- Outlier level start (end of low slope ramp, start of high slope ramp in different color, 'C' in Figure 2)

We have focused on localizing the primary level. For the secondary level, the outlier level and the primary level standard deviation, we have not yet revealed any correlation. Therefore, we suggest to use a fixed ratio to the primary level for these levels. We suggest to set the secondary level at half the primary level, the outlier level at double the primary level, and the primary standard deviation to $\pm 10\%$ around the primary level.

To establish a relationship between that density level and the measures, the primary level was manually set by an expert to give a good visualization of the MET volumes. We have investigated direct correlation between the manual levels and the measures, i.e., performed analysis of the different measures separately.

In Figure 2, the density histogram, the mean gradient and the mean curvature are shown for a TMV volume. The number of bins for calculating the different functions is set to 256, which is higher than the accuracy of the extracted estimates.

A central point to investigate in these plots is whether they are based on multiple distributions, which could arise from differences between spurious blobs in the volumes and shapes of actual molecules. In the following subsections the measures are described. The results are shown and discussed in Section 4.

The density histogram is the basis for measure 1. This is a "standard" histogram that shows the number of voxels that fall within a certain density bin. We investigate how the expert levels are distributed in the histogram. The ideal, but a little boring, result would be that the levels correspond to a single histogram percentile. This would mean that it would suffice to visualize a fixed fraction of a MET volume.

The number of components with size filtering is the basis for measure 2. This creates a function over the density values as for the histogram, representing the number of components at one density bin within the specified size range. To obtain this number, the lower bin value is used as threshold and all voxels with a value equal to or higher than this should belong to a component. If two voxels are 26-connected, then they belong to the same component. The size filter is applied to increase the "hit rate" in the density region of most interest. To extract the primary level estimate from this function, the position of the maximum is simply used. The size estimates have been manually calculated from the object diameters and their basic shape (round or tubular), with some margin.

The mean gradient histogram is the basis for measure 3. A gradient magnitude image is calculated using the first order derivative of a Gaussian kernel with a sigma related to the approximate diameter in voxels of the components of interest in the images. The gradient magnitude for each voxel is

$$|\frac{\partial f}{\partial x}| + |\frac{\partial f}{\partial y}| + |\frac{\partial f}{\partial z}| \tag{1}$$

where f is the density. The mean gradient magnitude value for the set of voxels within the density bin is then calculated. From the plots of this function, we have identified a reoccurring plateau starting approximately around the primary level chosen by the expert. We have defined the starting point for this plateau as the first local minimum of a smoothed derivative of the mean gradient measure. The smoothing has been done with a Gaussian filter with sigma set to 5, half the size of one subpart of the IgG molecule.

The mean curvature histogram is the basis for measure 4. It follows the curvature approach used by Kindlmann [4]. Essentially, these values measure the mean curvature for voxels at a particular intensity level. A value of zero would mean that the voxels within the corresponding intensity interval is not a part of an isosurface that has a strong curvature.

The first and second order partial derivatives needed for the gradient $\mathbf{g}$ and the Hessian matrix $\mathbf{H}$ are calculated using combinations of first and second order derivatives of the Gaussian kernel. Then the surface normal

$$\mathbf{n} = \mathbf{g}/|\mathbf{g}| \tag{2}$$

and the projection matrix

$$\mathbf{P} = \mathbf{I} - \mathbf{n}\mathbf{n}^T \tag{3}$$

where $\mathbf{I}$ is the identity matrix, are calculated. The matrix $\mathbf{P}$ projects onto the tangent plane of the isosurface. Next, the matrix

$$\mathbf{G} = \text{-}\mathbf{PHP}/|\mathbf{g}| \tag{4}$$

referred to by Kindlmann as the geometry tensor, is formed and from that the trace T and Frobenius norm F. The mean curvature is calculated as:

$$\kappa_1 = (T + \sqrt{F^2 - T^2})/2 \tag{5}$$

$$\kappa_2 = (T - \sqrt{2F^2 - T^2})/2 \tag{6}$$

Then, the mean curvature for a single voxel is $(\kappa_1 + \kappa_2)/2$. As in the mean gradient calculation, the total mean curvature for one density bin is calculated to acquire the measure. To match it with the primary level,

Dataset	Sigma	Size threshold (voxels)
RNAP II	2.0	2500
IgG	1.0	500
CEACAM1	1.0	500
TMV	2.0	50000

Table 2: Parameters for Gaussian and size filtering

we have identified a plateau for this measure as well. We have defined its starting point in the same way as for the gradient measure, i.e., as the first local minimum of a smoothed derivative of the mean curvature measure.

A priori information has been used when calculating the gradient and curvature measures, in the form of setting sigma to 1/10 of an estimation of the smallest component diameter. This will preserve the main structure of the components, but remove some of the noise.

For volume rendering, pre-integrated ray casting was used. Ray casting provides the possibility to visualize more information than an isosurface rendering. We claim that the user can get a better feeling for the data in the volume, by for example also taking densities around an isolevel into consideration. In order to still have an exact visualization with transfer functions which can contain high frequency changes, the pre-integration step is necessary. The methods were implemented in C++ partly using routines from the National Library of Medicine Insight Segmentation and Registration Toolkit (ITK) [3].

4 RESULTS AND DISCUSSION

We have not found any clearly multimodal distribution in any of the feature functions. From visual inspection of the density histogram, the distribution is close to a gamma distribution regardless of what kind of molecules there are in the solution. For the other three feature measures, there is more variation. In Figure 9, the correlation between the expert levels and the measures are shown. The used manually set filter parameters are given in Table 2. For the immediate visualization of a volume when it is opened, it should not be required to enter such information, but preliminary tests show that using a standard value for sigma will still give feasible results.

Measure 1 was visually evaluated, whereas for measures 2-4, a performance index has also been calculated for each of the measures. First a line $y = kx + m$ has been fitted to the data using a least square error norm. The performance index is calculated as

$$P = 100\frac{k}{\text{error}} \qquad (7)$$

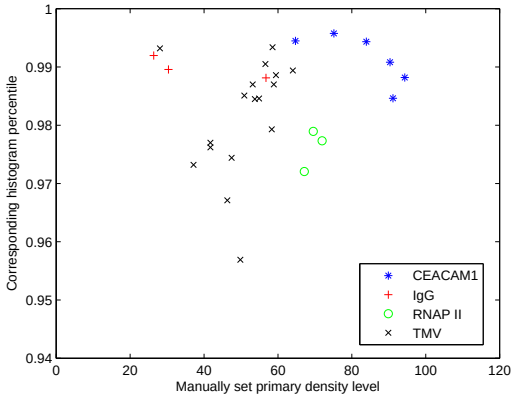

Figure 3: Primary levels chosen by an expert for the 29 investigated volumes plotted against their corresponding density histogram percentiles.

This index evaluates the discriminative power of a measure, but it is nothing more than an inverted correlation to the error, which is multiplied with k to remove the effect of the value range of the measure. The error is calculated as the mean of absolute errors.

1. Density histogram, Figure 3: The aim is to find a decorrelation when calculating the percentile value from the manual level. That is, the manual levels would optimally correspond to a single percentile level, forming a horizontal line. In Figure 3, a decorrelation tendency can be seen as not all the data seem to be spread around any diagonal "correlation line". Another observation is that the primary level is always above 95% for the investigated data sets. A comparison of an expert visualization and a visualization using the 99-percentile of the density histogram as the primary level is shown for RNA Polymerase II in Figures 4 and 5, respectively. The same comparison is shown for TMV in Figures 6 and 7, respectively.

2. Connected components, Figure 9: Considering all data sets, there is only some correlation. When excluding the sets with the largest molecules, the TMV data sets and the RNA Polymerase II data sets, a linear tendency can be seen. The performance index is 4.1 using all data sets.

3. Gradient, Figure 9: This measure exhibits the highest correlation tendency to the expert primary level, with a performance index of 7.5. This indicates that the molecules of interest have a more homogeneous internal density structure than the lower intensity artefacts.

4. Curvature, Figure 9: For the curvature measure, there seems to be a weak linear tendency. The performance index is 3.0. Hence, the difference in isosurface curvature between the biological molecules and other structures does not seem very significant.

Combining the measures using their performance indices as averaging weights gives the result shown in Figure 8. The performance index of this combined measure is 7.4, which is lower than for the gradient based measure.

When testing on an Intel E5430 2.66GHz CPU, loading and processing a 256^3 volume took approximately one minute. In an interactive application, the generation of a transfer function needs to be faster, a few seconds would be preferable, which should be feasible with optimization.

5 CONCLUSION

Four measures have been investigated regarding their potential for automatically generating a first visualization of MET volume data. We see all as interesting measures in this context, but the gradient based measure stand out as giving the best estimate of the primary level. We therefore suggest to use gradient based analysis for best accuracy when setting the primary level. Another simple but interesting result is that the primary level of interest for the investigated data sets is always in the top 5% of the volumes, in terms of density. Since this percentile measure is fast to compute, it is a good basic measure for instant automatic visualization, especially of large MET volumes. It could also be used as a control measure when calculating the primary level in a more exact way.

It is suitable that the gradient measure with the highest performance index also is the second easiest measure to calculate, after the histogram percentile, although it still takes around 15 seconds for a typical volume to be processed for this measure. In terms of algorithmic complexity, the gradient measure is based on separable filtering, so it will scale nicely for larger volumes.

Our next step is to step up a scale in terms of feature calculation, to make the distinction of objects of interest and other structures easier. One path would be to use region growing methods and explore different components using suitable shape descriptors.

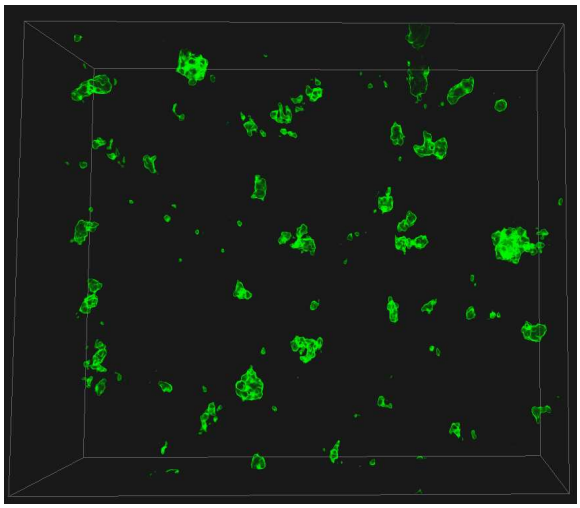

Figure 4: Visualization of a RNA Polymerase II volume, with an expert set primary level.

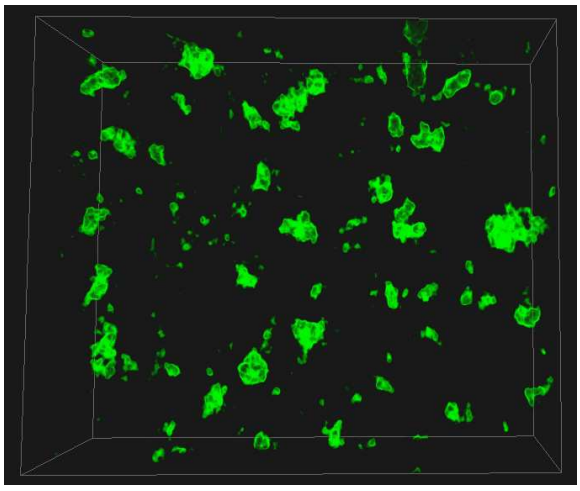

Figure 5: Visualization of the same RNA Polymerase II volume as in Figure 4, setting the primary level at the 99-percentile of the density histogram, but also showing $\pm$0.8% around that level.

ACKNOWLEDGEMENTS

This work is funded through the *Visualization Program* by Knowledge Foundation, Vårdal Foundation, Foundation for Strategic Research, VINNOVA, and Invest in Sweden Agency.

We would like to thank Roger Kornberg Laboratory, Dept. of Structural Biology at Stanford University Medical School, USA; Elenor Hauzenberger and Lars-Göran Öfverstedt at Sidec AB, Sweden; Sara Sandin, et al. [9]; Esther Klaile, et al [5], for providing the data and to Daniel Evestedt, SenseGraphics AB, for partly implementing the pre-integrated ray caster. We also thank the anonymous reviewers for their constructive criticism.

REFERENCES

[1] Chandrajit L. Bajaj, Valerio Pascucci, and Daniel R. Schikore. The contour spectrum. In *Proc. Vis. '97*, pages 167–173, 1997.

[2] Klaus Engel, Martin Kraus, and Thomas Ertl. High-quality pre-integrated volume rendering using hardware-accelerated pixel shading. In *SIGGRAPH/Eurographics Workshop on Graphics Hardware*, pages 9–16. ACM, 2001.

[3] Luis Ibanez, William Schroeder, Lydia Ng, and Josh Cates. *The ITK Software Guide*. 2005.

[4] Gordon Kindlmann, Ross Whitaker, Tolga Tasdizen, and Torsten Möller. Curvature-based transfer functions for direct volume rendering:

methods and applications. In *Proceedings Vis. '03*, page 67, 2003.

[5] Esther Klaile, Olga Vorontsova, Kristmundur Sigmundsson, Mario M. Müller, Bernhard B. Singer, Lars-Göran Öfverstedt, Stina Svensson, Ulf Skoglund, and Björn Öbrink. The CEACAM1 N-terminal Ig domain mediates cis- and trans-binding and is essential for allosteric rearrangements of CEACAM1 microclusters. *Journal of Cell Biology*, 187(4):553–567, 2009.

[6] Joe Kniss, Simon Premoze, Milan Ikits, Aaron Lefohn, Charles Hansen, and Emil Praun. Gaussian transfer functions for multi-field volume visualization. In *Proc. Vis. '03*, page 65, 2003.

[7] Hanspeter Pfister, Bill Lorensen, Chandrajit Bajaj, Gordon Kindlmann, Will Schroeder, Lisa Sobierajski Avila, Ken Martin, Raghu Machiraju, and Jinho Lee. The transfer function bake-off. *IEEE Comput. Graph. Appl.*, 21(3):16–22, 2001.

[8] Christof Rezk Salama, Maik Keller, and Peter Kohlmann. High-level user interfaces for transfer function design with semantics. *IEEE TVCG*, 12(5):1021–1028, 2006.

[9] Sara Sandin, Lars-Göran Öfverstedt, Ann-Charlotte Wikström, Örjan Wrange, and Ulf Skoglund. Structure and flexibility of individual immunoglobulin G molecules in solution. *Structure*, 12(3):409–415, 2004.

[10] Ulf Skoglund, Lars-Göran Öfverstedt, Roger M. Burnett, and Gérard Bricogne. Maximum-entropy three-dimensional reconstruction with deconvolution of the contrast transfer function: A test application with adenovirus. *Journal of Structural Biology*, 117:173–188, 1996.

[11] Gunther H. Weber, Scott E. Dillard, Hamish Carr, Valerio Pascucci, and Bernd Hamann. Topology-controlled volume rendering. *IEEE TVCG*, 13(2):330–341, 2007.

[12] Jianlong Zhou and Masahiro Takatsuka. Automatic transfer function generation using contour tree controlled residue flow model and color harmonics. *IEEE TVCG*, 15(6):1481–1488, 2009.

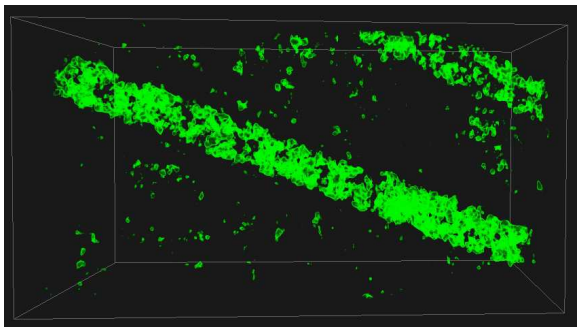

Figure 6: Visualization of a MET volume with a Tobacco Mosaic Virus (TMV) using an expert set primary level.

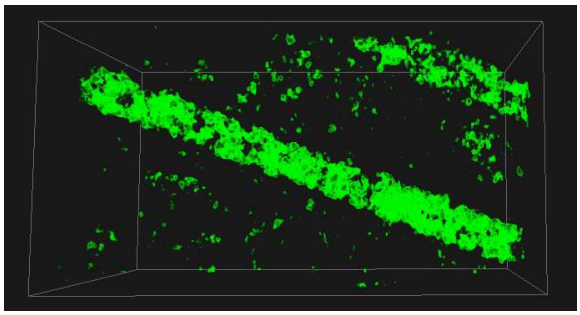

Figure 7: Visualization of the same MET volume of a TMV as in Figure 6 setting the primary level at the 99-percentile of the density histogram, but also showing $\pm$0.8% around that level.

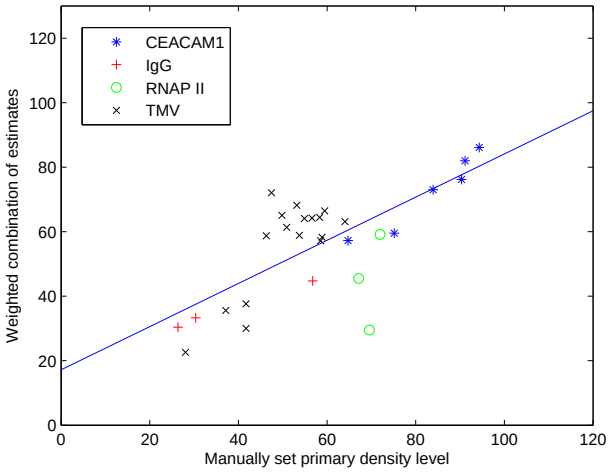

Figure 8: How a weighted combination of the measures correlates to the expert set primary level. The fitted line is used to calculate the performance index of the combined measure.

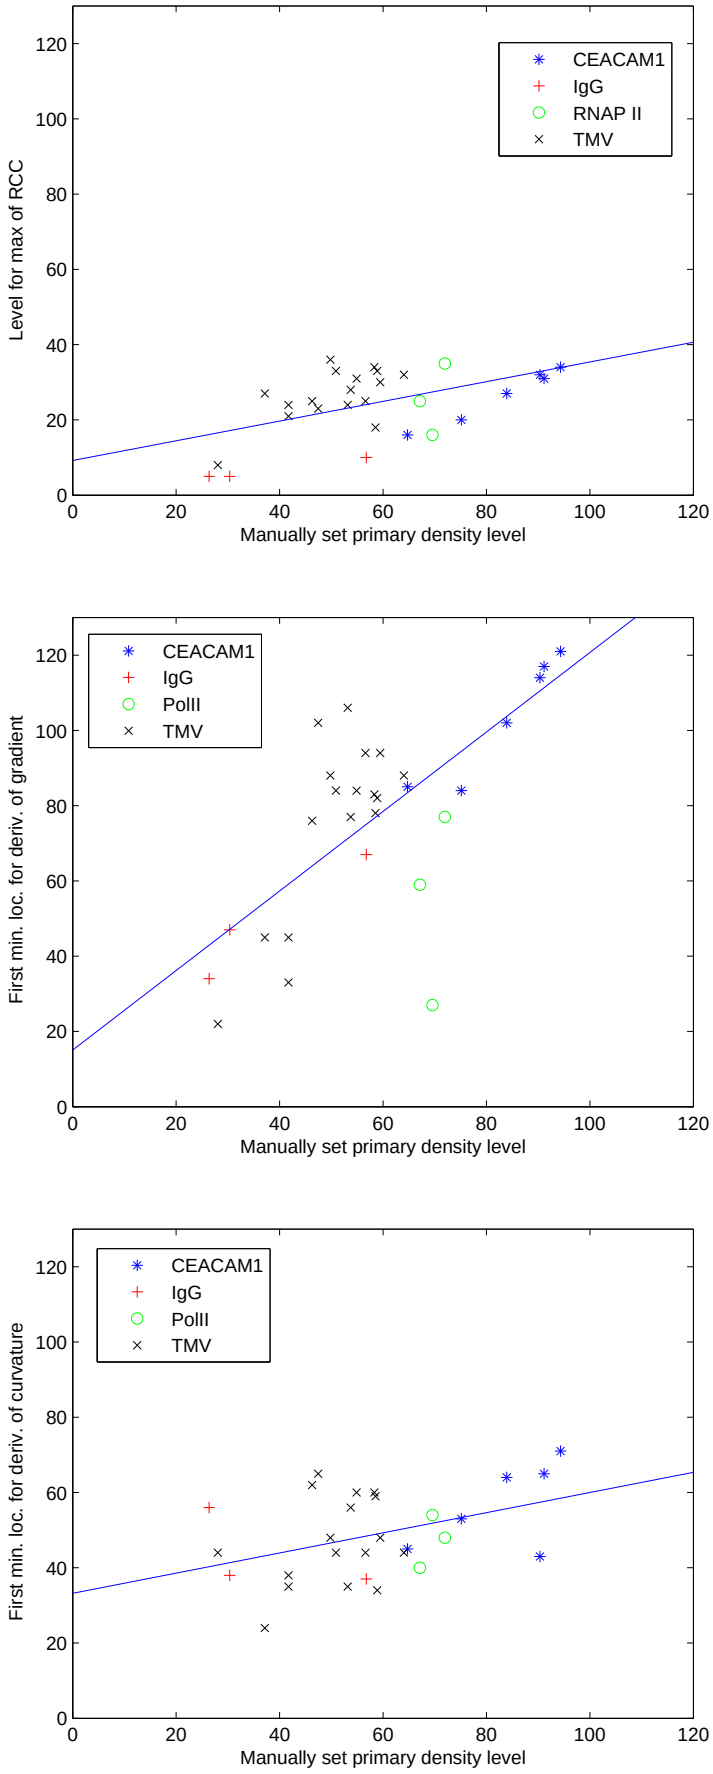

Figure 9: How the expert set primary level correlates to the extracted measures. The fitted line is used for calculating the performance index of each measure. Top: For the connected component measure there is a weak correlation to the expert level. Middle: The gradient measure shows correlation. Bottom: The curvature measure shows a weak correlation.