Visualization of 3D Cluster Results for Medical Tomographic Image Data

Sylvia Glaßer, Kai Lawonn and Bernhard Preim
Department of Simulation and Graphics, Otto-von-Guericke University, Magdeburg, Germany

Keywords: Computer Graphics, I.3.5, Picture/Image Generation, Display algorithms.

Abstract: We present an approach for the 3D visualization of clustered tomographic image data using the example of breast perfusion image data. Our visualization provides fast visual access to the amount of clusters, cluster size, presence and amount of outliers, and the spatial extent as well as the spatial orientation of the clusters. The spatial perception of a cluster's elements is improved with a connection via geometric primitives and appropriate shading styles and color mapping. Our technique can be easily adapted to any cluster result arising from medical tomographic image data.

1 INTRODUCTION

Medical image data sets become larger and more complex, thus increasing the demand for computer support during the evaluation by a biomedical expert. Hence, clustering is important to group objects with similar attributes, e.g., voxels representing certain tissue types. The appropriateness of a cluster algorithm depends on the medical problem and image parameters. To improve the adoption of a clustering to a specific problem, the biomedical expert has to evaluate the quality of the clustering result, like the topology of the clusters and the spatial orientation of the structure that has been decomposed.

We present a new 3D clustering view to visualize clustering results of medical image data, e.g., 4D temporal perfusion data. For the evaluation of perfusion data sets, clinical diagnosis aims at the identification of areas that exhibit similar characteristics. For example, clustering of breast tumor perfusion data is employed to evaluate the tumor's heterogeneity (Preim et al., 2012). Beyond grouping of voxels with similar perfusion characteristics, clustering is applied to various medical image data, such as fMRI data. Hence, peaks in a histogram often represent overlaps of tissue types that are not particularly interesting. To detect the interesting data parts, clustering is often employed to analyze the gradient versus the intensity histogram space (Maciejewski et al., 2009). Furthermore, cluster analysis is carried out to generate transfer functions for this kind of data (Maciejewski et al., 2013). In addition, cluster analysis is also applied to identify regions with similar boundary characteristics, which in turn are used for transfer function specification (Sereda et al., 2006). DTI data forms another important medical application area for clustering, where fiber tracts are clustered (Moberts et al., 2005).

We illustrate our new 3D cluster view based on the example of clustering of medical perfusion MRI data for breast cancer diagnosis. Perfusion MRI exhibits a higher sensitivity than conventional X-ray mammography in younger women and is often consulted to confirm benignity (i.e., non-cancer) or malignancy (i.e., cancer) of a tumor. Contrast-enhanced perfusion MRI data is acquired to analyze the contrast agent kinetics of the tumor. A breast tumor is considered as malignant as its most malignant tumor part.

Consequently, biomedical researchers try to detect new correlations between heterogeneity and malignancy based on the clustering result of breast tumors. For the evaluation of a clustering's quality, the cluster shapes and spatial orientations as well as their average enhancement are assessed. Also, the topology and the connectivity of the clusters is evaluated such that conclusions are made about a necrosis in the tumor center or varying enhancement kinetics at the tumor boundary. Generally, a standard 2D slice view is analyzed since this visualization does not suffer from any occlusion. However, we aim at a 3D representation to not only show the spatial orientation of the clusters, but also how they penetrate each other.

When it comes to a 3D visualization of different groups of data (i.e., different groups of voxels that could be arbitrarily shaped and be nested into each

Glaßer S., Lawonn K. and Preim B..
Visualization of 3D Cluster Results for Medical Tomographic Image Data.
DOI: 10.5220/0004719901690176
In *Proceedings of the 9th International Conference on Computer Graphics Theory and Applications* (GRAPP-2014), pages 169-176
ISBN: 978-989-758-002-4

other), in general each group is visualized with its boundary surface and one is directed to the embedded surface problem. We avoid this by employing a 3D scatter plot-like view. The objects (in our case the voxels) are mapped to small spheres in the 3D scatter plot-like coordinate system. Thus, we can avoid occlusion, since no surfaces or blocks of data are rendered. To emphasize the connection of the data points to the corresponding clusters, we employ tubes and use color-coding. Additional information about the clusters, like average relative enhancement curve, is provided with a necklace map legend (Speckmann and Verbeek, 2010). Our framework is developed based on the clustering of medical perfusion data. However, our method is feasible for each cluster 3D visualization arising from other medical areas.

In summary, we present a visualization with the following contributions. First, we describe a center point-oriented projection of a tumor into a sphere model, where different spheres represent different neighbor layers. Now, each voxel can be rendered as small object on such a sphere's surface. Second, we introduce a 3D scatter plot-like visualization for breast clustering results employing the sphere model and featuring well-fitted tubes. The visualization comprises selected shading styles to enhance the different scatter plot parts and a perception-based color mapping to enhance visual differentiation between clusters. Furthermore, it is adapted to the extent of the tumor. Third, a necklace map-based legend provides different information about the clusters. We evaluate our visualization in a quality user study.

2 RELATED WORK

To improve and support the radiologist's work, advanced methods including visual analytic solutions were developed. Lundström and Persson characterized the visual analytic tasks in diagnostic imaging (Lundström and Persson, 2011). Hence, the determination of shape, size, and relative position of different parts of the anatomy was identified as important component of a radiologist's image review work, albeit it was considered less important than the diagnosis of primary and secondary findings in the image data in an efficient way. For the evaluation, Glaßer et al. stated that density-based clustering techniques are well suited for the underlying breast perfusion image data due to the arbitrarily shaped clusters as well as outlier detection (Glaßer et al., 2013). The aim of the current work is not to identify the best clustering but to visualize an extracted cluster result such that important information, e.g., number of clusters, number of outliers, can be extracted. A Multifield data visualization can be carried out with geometric presentations like the scatter plot (Wong and Bergeron, 1997) for 2D object spaces or the parallel coordinates plot (Inselberg and Dimsdale, 1990) for higher dimensional object spaces without spatial information. To maintain spatial information in high-dimensional cluster visualization, Linsen et al. presented a surface extraction approach for spatial visualization of multifield clustered data (Linsen et al., 2008). Hence, boundary surfaces are extracted and semi-transparently visualized in a 3D star coordinate layout to reveal nested clusters. Poco and colleagues employed a least square projection to map high-dimensional data to 3D visual spaces (Poco et al., 2011). Their framework enables the mapping of multidimensional non-spatial data as well as the feature space of multi-variate spatial data. Both approaches employ surface representations for their cluster results. Due to the limited number of voxels in a medical data set, we do not want to visualize surfaces but rather present each object directly.

In (Zhang et al., 2006), a method to visualize gene clusters in 3D is described. First, a spring model is used to locate genes within a cluster into InfoCubes. Afterwards, the same method is used to allocate the InfoCubes into 3D space. The algorithm avoids the space partition problem. Unfortunately, clusters on our tumor data sets can be entwined around each other such that we need a representation of such cases. The approach by Quigley deals with large data sets and generates hierarchical compound graphs (Quigley, 2001). The clusters are represented as a graph where the nodes act as the clusters. Yang et al. presented a system for the processing of DNA microarray data (Yang et al., 2003). They introduce a space-undivided and a space-divided 3D gene plot. In the first plot, every gene is assigned to a sphere and the cluster membership is color-coded. The second plot splits the 3D space into cubes and puts the clusters into these cubes. In the first case, clusters seem to be huge point clouds, and in the second case, no spatial information about the clusters exists. Aono and Kabayashi developed a visual interface for non-technical users to understand the output from clustering algorithms (Aono and Kobayashi, 2011). Their algorithm displays clusters which are projected on a 3D subspace based on the user's keyword input. Thus, the computed recommendations for the 3D visualization cannot present the spatial impression on the tumor.

In contrast, our approach aims at direct mapping of each voxel to a data point in the 3D space, which is possible due to our small data set extents. Similar to our work, Kosara et al. presented the VoxelPlot, a 3D

scatter plot that interactively links scientific and information visualization (Kosara et al., 2004). A voxel-based representation of clustered / grouped data for perfusion MRI was presented in (Oeltze et al., 2007). Sanftmann and Weiskopf introduced the illuminated scatter plot, where shape perception is improved by applying a dedicated illumination technique to the 3D scatter plot point cloud representation (Sanftmann and Weiskopf, 2009). Piringer and colleagues map the distance to color and the point size of the scatter plots (Piringer et al., 2004). In contrast to these approaches, we do not only employ color or shading to represent the type of cluster, but also add geometric objects into our 3D scatter plot to enhance the connections of cluster voxels to cluster centers. This is similar to the work presented in (Healey et al., 2001), where contextual cues in the final 3D scatter plot-based visualizations for large, multidimensional collections of data are included. Another example related to the aspect of connected cluster elements is the butterfly plot (Schreck et al., 2008). Hence, multi-dimensional point clouds are projected into 2D space and clusters are visualized by compact shapes that enclose all members of a given point cloud.

Our visualization method is extended by a legend that is based on the necklace map (Speckmann and Verbeek, 2010). This technique is related to boundary maps, which map labels around the map in the center, and is well suited for cluster visualization.

3 MATERIAL AND METHODS

In this section, we describe the steps to create the 3D clustering view, see Figure 1. At step I, the neighborhood model is extracted (described in Sec. 3.1). At step II, the voxel size and tumor extent is determined (see Sec. 3.2). At step III, we project the voxels onto spheres (refer to Sec. 3.3), connect the cluster elements with their associated cluster center points via tubes (presented in Sec. 3.4), and reshape the spheres according to the tumor's extent. A legend that is based on the necklace map is added to the 3D visualization to provide additional information about the clusters, described in Section 3.5. The assignment of color and shading styles is presented in the last part of this section.

3.1 Neighborhood Model for Voxel Ordering

Given a voxel in a 3D tomographic data set, the corresponding neighbors are defined as adjacent voxels that share

- a face (yielding 6 neighbors),
- an edge (yielding 18 neighbors),
- or a vertex (yielding 26 neighbors).

We choose the edge-based neighborhood relationship. We build a neighborhood structure by an iterative process. First, we start at an initial voxel v_{init} that is typically the centroid of the tumor $\mathscr{T}$. If the centroid is not part of the tumor due to an irregular tumor shape, we choose its nearest neighbor contained in the tumor. The first layer $\mathscr{L}_1$ of our neighborhood structure contains only v_{init}. The second layer $\mathscr{L}_2$ includes the neighbors $v_i \in \mathscr{T}$ of v_{init} that are also part of the tumor. For simplicity, we denote $\text{Neigh}(v)$ as the neighbors of v in $\mathscr{T}$. The i-th layer can be written as:

$$\mathscr{L}_i = \bigcup_{v \in \mathscr{L}_{i-1}} \text{Neigh}(v) \setminus \mathscr{L}_{i-1}. \quad (1)$$

3.2 Integration of Voxel and Tumor Extent

The voxel extent is determined by the image matrix that defines the image plane resolution, i.e., the voxel's width and height. The voxel's depth is set down by the slice thickness, which usually differs from the image plane resolution. We rescale the voxel's depth v_z such that the voxel's width and height equals uniform length.

The tumor extent is characterized with a principal component analysis. The voxel's x, y and z indices serve as input, whereas z was rescaled with v_z. Hence, the principal components match the eigenvectors of the covariance matrix that was constructed for the principal component analysis. With the three eigenvalues corresponding to the three principal components, we obtain the tumor's three main extents.

3.3 Projection onto a Sphere

After we extracted the neighborhood structure, we build up our cluster visualization. Each layer of neighborhood information is projected onto a sphere. Therefore, we translate a sphere of radius $r = i^2$ to the initial voxel, i.e., the centroid. Afterwards, we determine the intersection point of the sphere with the line constructed by the initial voxel and an element in the i-th layer $\mathscr{L}_i$. In detail, we translate the voxel such that the initial voxel is in the origin. Given the midpoint $\mathbf{q}_j^i$ of $v_j \in \mathscr{L}_i$, we obtain the intersection point $\mathbf{p}_j^i$ by calculating

$$\mathbf{p}_j^i = \mathbf{q}_j^i \sqrt{\frac{r^2}{(\mathbf{q}_j^i)^2}}. \quad (2)$$

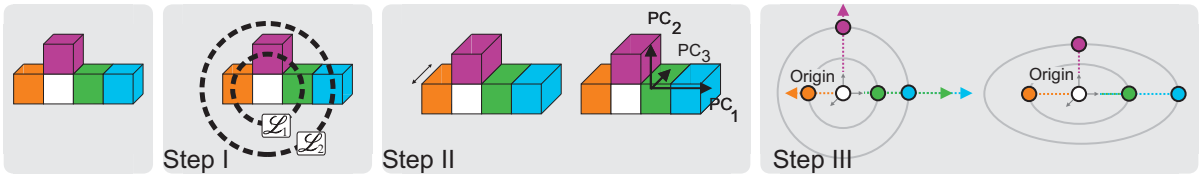

Figure 1: Scheme of our cluster visualization pipeline. The example holds five voxels, each forming its own cluster. In step I, we create a sphere model, holding all voxels of the tumor. In step II, the voxel dimensions are adapted to the image resolution. A principal component analysis yields the main directions of the tumor. In step III, the 3D clustering visualization is generated by positioning the voxels onto spheres. The sphere-shaped visualization is adapted to the main tumor extents, approximated with the principal components.

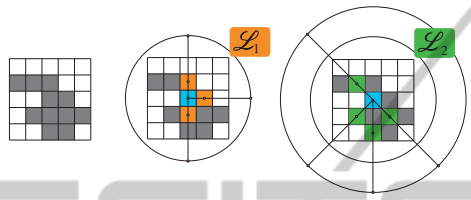

Figure 2: We start with the initial voxel (cyan color-coded), see (a). Next, we determine $\mathscr{L}_1$ and calculate the intersection point of the line constructed by the initial voxel and an element of $\mathscr{L}_1$ with the circle, see (b). In (c), we obtain the intersection points of the second layer $\mathscr{L}_2$.

See Figure 2 for an illustration in 2D.

3.4 Connection of Cluster Elements and their Associated Midpoints

To support the perception of clustered elements, we connect all elements and their associated midpoints. The connections, shaped like tubes, are motivated by the natural model of neurons and their dendrites as well as blobby surfaces (Blinn, 1982). We employ a visualization which connects the start points, i.e., the extracted positions for each voxel, with the end point, i.e., the cluster center, by smooth organic linkage. Hence, the algorithm comprises three steps. First, we connect the start point and the end point with a straight line. As an optional step, we add several equidistant points on the line. Afterwards, we use the normalized vector which points to the start point. We scale this vector to one tenth of the length from the end point to the start point and add it to the added points on the line. The scaling linearly changes to zero such that the first points are more translated than the last ones. This approach yields a bending of the tubes and reduced visual clutter. The second step is about generating a cubical spline to connect the end point and the start point via the middle points. The last step generates a Frenet frame around the curve. Afterwards, we use this frame to generate a tube around the curve. Inspired by the following formula:

Figure 3: An example of a visually pleasing tube connecting the start point to the end point.

$$\text{thick}(x) = \text{radius}\left(1 - \left(\frac{x}{\text{d}}\right)^2\right)^2, \qquad (3)$$

we generate a function which determines the thickness of the tube at every point. Let r_1 be the radius of the sphere at the start position, r_2 the radius of the sphere at the end point, and d denotes the distance between the start and the end position. Additionally, we define $r' = 0.25\min(r_1, r_2)$. Furthermore, $r_a = r_1 - r'$ and $r_b = r_2 - r'$. Then, the $\text{thick}(x)$ function with $x \in [0, \text{d}]$ can be written as:

$$\text{thick}(x) = \begin{cases} r_a \cdot \left(1 - \left(\frac{4x}{\text{d}}\right)^2\right)^2 + r' & \text{if } x \leq \frac{\text{d}}{4} \\ r_b \cdot \left(1 - \left(\frac{4(\text{d}-x)}{\text{d}}\right)^2\right)^2 + r' & \text{if } x \geq \frac{3}{4}\text{d} \\ r' & \text{otherwise.} \end{cases} \qquad (4)$$

The construction of $\text{thick}(x)$ ensures a smooth change from r_1 for $x = 0$ to r_2 for $x = d$. First, the value $\text{thick}(x)$ decreases smoothly from r_1 to r' for $x \in [0, \text{d}/4]$ and keeps its value r' for $x \in [\text{d}/4, 3\text{d}/4]$. For $x > 3\text{d}/4$ the function is smoothly increased to the value r_2. See Figure 3 for an example.

3.5 Combination with a Necklace Map

To assess a clustering's quality, additional information to the cluster's spatial orientation should be provided. These information may include the cluster's

average value of a parameter, its standard deviation, or its size, etc.

For perfusion imaging, the average perfusion enhancement curve is of major interest from the radiologist's point of view. For each cluster, we provide its average curve in a legend that is based on a necklace map (Speckmann and Verbeek, 2010). The necklace map was developed for 2D maps and is similar to cartograms or choropleth maps. However, it arranges symbols (e.g., circles) around the initial map in a linear ordering. Thus, the symbols can carry information but no occlusion arises.

The necklace map is perfectly suited for our method, since we want to display additional information for a group of well-defined objects – our clusters – albeit we want to accomplish a 3D scene instead of a 2D scene. The according symbol for each cluster is a circle that is color-coded with the corresponding cluster color. All circles are arranged on the necklace ordered by cluster size, starting in the top right. The necklace itself is an ellipse that is obtained by scaling a circle with tumor extents extracted in Section 3.2. The necklace legend is depicted in Figure 4. The size of the necklace pearls is linearly decreased, and not proportional to the cluster size due to strong variations of cluster sizes, e.g., a cluster may contain 3 or 500 voxels.

3.6 Representation of Clusters

The tubes are rendered with a Phong shading to improve the shape perception. To support the visualization of the start point, and thus a good differentiation from the Phong-shaded connecting tube, we apply Fresnel shading to the start point.

We assign a specific color to each cluster and its start points. For the color assignment, we employ the CIELAB color space to establish high perceptual color contrast. The first cluster is always assigned to orange, with the corresponding CIELAB components $L = 67$, $a = 43$, and $b = 74$. For the n remaining clusters, we extract n colors with the following routine. We place a circle in the CIELAB space. The circle's center is set to $(a = 0; b = 0)$. Given the a- and b-value for orange, we define its radius r with $r^2 = a^2 + b^2$. Then, we define the angle $\alpha := \arcsin(b/r)$. Next, we compute new values for a and b by increasing α with β, where β ranges from 0 to $2\pi/n$ in n steps. We now compute $a = r \cdot \cos(\alpha + \beta)$ and $b = r \cdot \sin(\alpha + \beta)$, and combine them with the starting value for L.

For the final representation, we scale the spheres holding the voxels according to the tumor extent to highlight the tumor's biological form. Thus, a tumor with a biological ellipsoid form yields a scaling of the spheres along the main tumor extents, whereas a sphere-shaped tumor will only cause minimal scaling.

4 RESULTS

We present our approach adapted to a breast tumor in Figure 4. There, an overview of our 3D scatter plot-like clustering visualization is provided. The user can examine the spatial extent of the single clusters and their spatial position. The spheres, which hold the different neighbor layers, are visualized with a Fresnel shading. We assign gray to these spheres. Outliers are presented with white colored spheres. Furthermore, a standard headlight is applied to the visualization.

The 3D view is accomplished with a necklace map. Thus, the user has a fast overview of all existing clusters and additional information are mapped. As demanded by our medical experts, we provide the cluster's average relative contrast agent enhancement as time-intensity curve. Our framework also holds a conventional slice view. To study a cluster's spatial orientation in more detail, the user can pick a cluster by clicking via the necklace map or the 2D view and study its extent in more detail. Now, Fresnel shading is interactively applied to the non-selected clusters with a gray color. This emphasizes the selected clusters and supports the individual examination, recall Figure 4.

In Figure 5, we loaded a brain tumor data set comprising a masked perfusion MRI data set. For brain perfusion, the parameter cerebral blood volume (CBV) is analyzed to detect hot spots, i.e., regions with elevated CBV values. The brain tumor was decomposed into five clusters with k-means based on the parameter CBV. Hence, no outliers are present. Instead of presenting the whole contrast agent enhancement curve, the pearls of the necklace map provide average CBV values and identify the pink cluster as cluster with highest averaged CBV value.

5 EVALUATION

We performed a qualitative evaluation of our cluster visualization method. Our goal was to assess the capability to express the topology of the clustering result in the 3D visualization. Hence, we focussed on the 3D representation of the clusters in combination with connecting tubes and presented a point-based scatter plot view and transparent isosurfaces for comparison, see Figure 6. The evaluation was conducted with one medical researcher and physician, and ten researchers

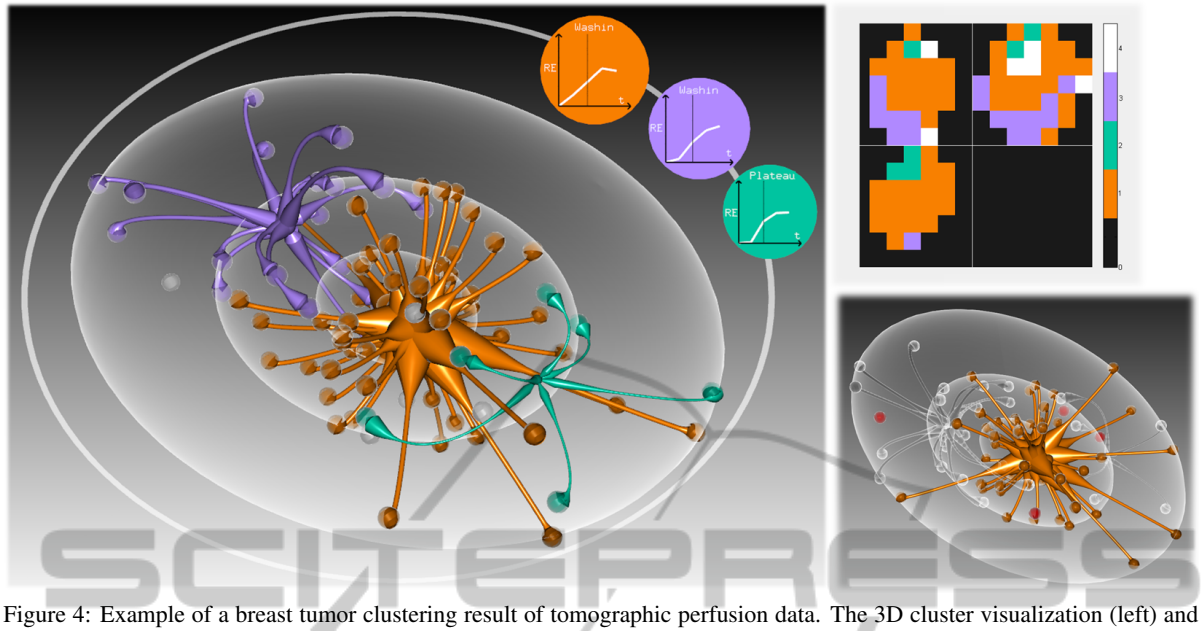

Figure 4: Example of a breast tumor clustering result of tomographic perfusion data. The 3D cluster visualization (left) and the 2D slice (top right) are integrated in our framework and reveal three clusters arranged at three layers. The necklace map (placed around the 3D visualization) and the slice view allow for a fast selection of clusters. Once the user selected a cluster, only this cluster is color-coded, and the remaining clusters are visualized with a Fresnel shading (bottom right). Outliers are highlighted with bright red after selection according to observations from our user study.

who are familiar with the visualization and evaluation of medical tomographic image data.

In this study, we asked the participants to handle a few minor tasks about

- the topology of the clustering result,
- the presence of outliers,
- and the impression about the tumor's boundary and shape.

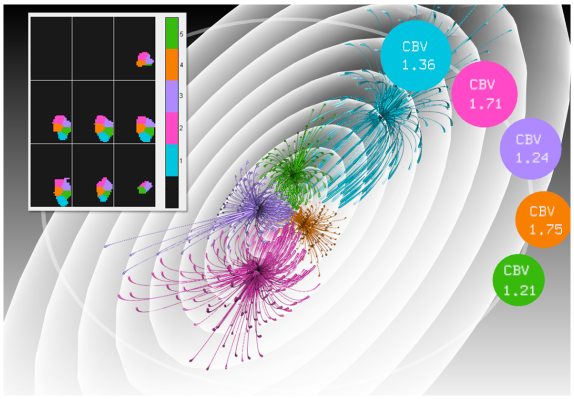

Figure 5: Illustration of the clustering result of a brain tumor. The tumor has an elongated shape and was clustered into five clusters, no outliers exist. The necklace map provides additional cluster information, i.e., the cluster's average parameter value of the cerebral blood volume (CBV). The slice view was added for a 2D view.

The topology of the clustering result was covered by the request to identify the largest cluster and to categorize the clustering result's topology in pre-defined categories. It was questioned if the participants could identify outliers and how many outliers there are with respect to number and percentage to the whole tumor. The last question concerned the tumor's boundary (from round to stellated boundary) and shape (from spherical to ellipsoidal). For each answer, we provided well-defined categories to make the answers of different participants comparable.

First, the participants were introduced to the three techniques by presenting a clustering result of a test data set (recall Fig. 6). Next, the subjects solved the tasks for nine examples, created by applying each visualization technique to three new data sets. The visualizations were presented such that no data set was consecutively shown, but after each presentation another data set with another 3D view was depicted. After answering the questionnaire, we asked the participants to rate the techniques according to their appropriateness to evaluate the clusters' topology, the required user interaction, as well as for additional feedback.

As a result, all volunteers correctly identified the largest cluster for each example. However, when it comes to the cluster number, our technique achieved better results than the other two methods, see Figure 7. This is very relevant, since a wrong number of

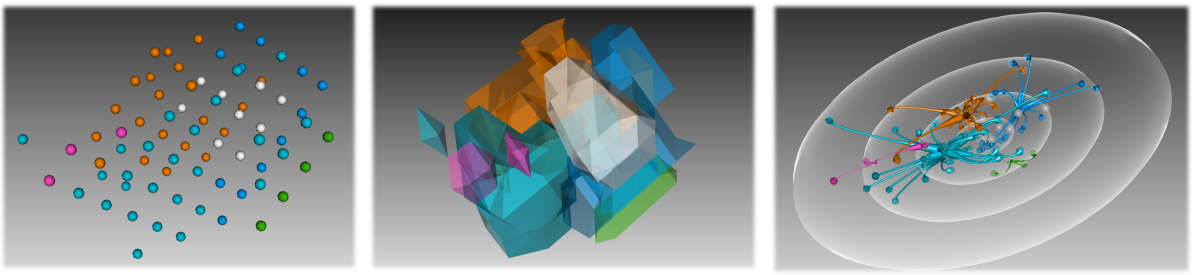

Figure 6: The three 3D views of a clustering result, including a point-based view (left), the isosurface view (center), and our approach (right).

clusters implies that some clusters were not detected at all or (very rare) that some single, spatially connected, clusters were interpreted as different clusters. The majority of the participants rated our method to be the most appropriate for evaluating the clusters' topology (8 out of 11).

Outliers were present in all examples, and with two exceptions (arising from our visualization and an isosurface-based view) for two single examples, all participants detected the outliers for each view. When it comes to tumor shape, the volunteers assign higher ratings, i.e., more stellated boundaries based on our visualization in comparison to the other two views. This is due to the arrangement of the voxels onto the spheres. Hence, our view suggests a more stellated boundary, which is a limitation. However, there was no trend with respect to the employed technique, when the participants should evaluate the tumor size via the number of voxels that were clustered. All participants had no difficulty with the interaction or with the visualization. But the majority (9 out of 11) needed less interaction with the 3D scene, e.g., camera rotation, with our method due to the good spatial impression of the connected clusters.

However, the study does not allow for a definitive statement. A further evaluation is required with more participants and data sets. In summary, the participants were able to fulfill the assigned tasks with our method very well. They rated it as best suited for the evaluation of a clustering topology. Furthermore, they preferred it due to the visual cluster connections via tubes and due to the lesser scene interaction in comparison to the other techniques. They asked for a more prominent color-coding of outliers, like bright red, which we included in the framework afterwards. It must also be stated that our visualization indicates a more stellated boundary due to the spatial representation of the voxels on the spheres. On the other hand, these representations reduce the amount of occlusion and thus the amount of required user interaction.

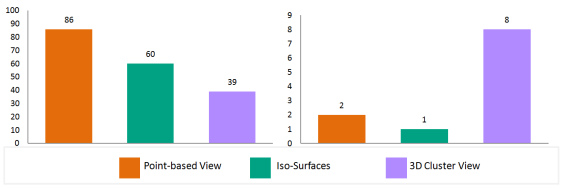

Figure 7: Bar diagrams illustrate the average squared error of the approximated number of clusters (left) by all users for the three techniques: point-based 3D view, isosurface view and our 3D scatter plot-like view. On the right, the number of users is presented that chose a technique as best suited for evaluation of the clustering's topology.

6 DISCUSSION AND FUTURE WORK

In this paper, we presented a novel method for 3D cluster visualization for analyzing medical tomographic image data. We applied our techniques to clustering results of breast perfusion data. However, our methods are feasible for each clustering arising from other medical areas. Our 3D scatter plot-like cluster visualization addresses the spatial information as well as the size of each cluster. A fast overview of how many clusters are present and how they are spatially aligned is presented. Geometric modeling enhances cluster connections. The 3D view is completed with a necklace map legend, maintaining that each cluster can be adressed. Hence, the user can select a cluster of interest for further exploration in the 3D view. While the appearance of the selected cluster does not change, color and shading of the remaining clusters change from opaque visibility to Fresnel.

For future work, we like to address an improved perception in case of a huge number of cluster elements. We think about a view-dependent transparency representation similar to the work by Günther and colleagues, who reduce the number of displayed lines by smoothly fading them out (Günther et al., 2011). Connections very close to the camera are opaque, whereas connections far away are rendered

transparent. This is done to improve the visual appearance of the tubes and to concentrate on the focus of the camera. Furthermore, we like to use the GPU to get a real-time interaction while changing the parameters of the cluster algorithm yielding an immediate result in the 3D cluster visualization after parameter changes.

ACKNOWLEDGEMENTS

This work was supported by the DFG Priority Program 1335: Scalable Visual Analytics.

REFERENCES

Aono, M. and Kobayashi, M. (2011). Text Document Cluster Analysis Through Visualization of 3D Projections. Technical report, IBM Research - Tokyo.

Blinn, J. F. (1982). A Generalization of Algebraic Surface Drawing. *ACM Trans. Graph.*, 1(3):235–256.

Glaßer, S., Niemann, U., Preim, B., and Spiliopoulou, M. (2013). Can we Distinguish Between Benign and Malignant Breast Tumors in DCE-MRI by Studying a Tumor's Most Suspect Region Only? In *Proc. of the IEEE Symposium on Computer-Based Medical Systems*, pages 59–64.

Günther, T., Bürger, K., Westermann, R., and Theisel, H. (2011). A View-Dependent and Inter-Frame Coherent Visualization of Integral Lines using Screen Contribution. In *Proc. of Vision, Modeling, and Visualization*, pages 215–222.

Healey, C., Stamant, R., and Chang, J. (2001). Assisted Visualization of E-Commerce Auction Agents. *Graphics Interface*, 1:201–208.

Inselberg, A. and Dimsdale, B. (1990). Parallel Coordinates: A Tool for Visualizing Multi-Dimensional Geometry. In *Proc. of the IEEE Conf. on Visualization*, pages 361–378.

Kosara, R., Sahling, G. N., and Hauser, H. (2004). Linking Scientific and Information Visualization with Interactive 3D Scatterplots. In *Proc. of WSCG Short Communication Papers*, pages 133–140.

Linsen, L., Van Long, T., Rosenthal, P., and Rosswog, S. (2008). Surface Extraction from Multi-Field Particle Volume Data using Multi-Dimensional Cluster Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1483–1490.

Lundström, C. and Persson, A. (2011). Characterizing Visual Analytics in Diagnostic Imaging. In *Proc. of the EuroVA International Workshop on Visual Analytics*, pages 1–4.

Maciejewski, R., Jang, Y., Woo, I., Jänicke, H., Gaither, K., and Ebert, D. (2013). Abstracting Attribute Space for Transfer Function Exploration and Design. *IEEE Transactions on Visualization and Computer Graphics*, 19(1):94–107.

Maciejewski, R., Woo, I., Chen, W., and Ebert, D. (2009). Structuring Feature Space: A Non-Parametric Method for Volumetric Transfer Function Generation. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1473–1480.

Moberts, B., Vilanova, A., and van Wijk, J. J. (2005). Evaluation of Fiber Clustering Methods for Diffusion Tensor Imaging. *IEEE Visualization*, pages 65–72.

Oeltze, S., Doleisch, H., Hauser, H., Muigg, P., and Preim, B. (2007). Interactive Visual Analysis of Perfusion Data. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1392–1399.

Piringer, H., Kosara, R., and Hauser, H. (2004). Interactive focus+ context visualization with linked 2d/3d scatterplots. In *Proc. of the Conference on Coordinated and Multiple Views in Exploratory Visualization*, pages 49–60. IEEE.

Poco, J., Etemadpour, R., Paulovich, F. V., Long, T. V., Rosenthal, P., Oliveira, M. C. F., Linsen, L., and Minghim, R. (2011). A Framework for Exploring Multidimensional Data with 3D Projections. *Computer Graphics Forum*, 30(3):1111–1120.

Preim, U., Glaßer, S., Preim, B., Fischbach, F., and Ricke, J. (2012). Computer-aided diagnosis in breast DCE-MRI – Quantification of the heterogeneity of breast lesions. *European Journal of Radiology*, 81(7):1532 – 1538.

Quigley, A. J. (2001). Large scale 3d clustering and abstraction. In *Selected papers from the Pan-Sydney workshop on Visualisation*, volume 2 of *VIP '00*, pages 117–118.

Sanftmann, H. and Weiskopf, D. (2009). Illuminated 3D Scatterplots. *Computer Graphics Forum*, 28(3):751–758.

Schreck, T., Schüßler, M., Zeilfelder, F., and Worm, K. (2008). Butterfly Plots for Visual Analysis of Large Point Cloud Data. In *Proc. of Conf. on Computer Graphics, Visualization and Computer Vision*, pages 33–40.

Sereda, P., Gerritsen, F. A., and Vilanova, A. (2006). Mirrored lh Histograms for the Visualization of Material Boundaries. In *Proc. of Vision, Modeling, and Visualization*, pages 237–244.

Speckmann, B. and Verbeek, K. (2010). Necklace Maps. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):881–889.

Wong, P. C. and Bergeron, R. D. (1997). 30 Years of Multidimensional Multivariate Visualization. In Nielson, G. M., H. H. and Muller, H., editors, *Scientific Visualization - Overviews, Methodologies, and Techniques*, pages 3–33. IEEE Computer Society Press.

Yang, Y., Chen, J. X., and Kim, W. (2003). Gene expression clustering and 3d visualization. *Computing in Science & Engineering*, 5(5):37–43.

Zhang, L., Sheng, W., and Liu, X. (2006). 3D Visualization of Gene Clusters. In *Proc. of Conf. in Computer Vision and Graphics*, pages 349–354.