Vision, Modeling, and Visualization (2013)
Michael Bronstein, Jean Favre, and Kai Hormann (Eds.)

Interactive Comparative Visualization of Multimodal Brain Tumor Segmentation Data

F. Lindemann[1], K. Laukamp[2], A. H. Jacobs[2], K. Hinrichs[1]

[1]Visualization & Computer Graphics Research Group, University of Münster, Germany
[2]European Institute for Molecular Imaging, Münster, Germany

Abstract

We present a visualization system for the analysis of multi-modal segmentation data of brain tumors. Our system is designed to allow researchers and doctors a further investigation of segmented tumor data beyond a quantitative assessment of size. This includes the efficient visual analysis of the shape and relative position of the different, often overlapping segmented data modalities, using high quality 3D renderings of the data. Furthermore, our system provides visualization methods to compare tumor segmentation volumes acquired at various points of time, which helps the user to explore changes in shape and size before and after treatment. We also employ two novel interactive diagrams which allow the user to quickly navigate and analyze overlapping tumor regions. All methods are assembled and linked in a multi-view framework.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Viewing algorithms, I.3.8 [Computer Graphics]: Applications—

1. Introduction

Gliomas are the most common primary tumors of the central nervous system with an incidence ranging between 6-8/100.000. Despite multi-modal therapeutic regimens (surgery, radiation, chemotherapy, anti-angiogenic therapy), the prognosis of high grade gliomas (World Health Organization Grade IV) is still poor with median survival being 1-2 years. In the past years, two major research areas have supported increased understanding of glioma biology: first, detailed molecular-genetic analysis revealing a step-wise accumulation of disease-specific genetic changes with some of them having prognostic relevance; second, development of multi-modal molecular and structural imaging based on positron emission tomography (PET) and magnetic resonance imaging (MRI). A combination of MRI and PET imaging is being used as primary diagnostic tool for tumor localization, grading, and delineation from functionally important neuronal tissue; thereafter, imaging is being used to quantify response to therapy, time of tumor recurrence and differentiation of biologically active tumor tissue from radiation necrosis [SPS*10]. One of the critical parameters influencing the disease's progression and outcome is the extent of the tumor in 3-dimensional space. The true tumor extent can be measured by a combination of various MRI- and PET-based parameters as they reveal complementary information on the tumor. The combined assessment of the tumor volume as measured by MRI and PET has direct impact on surgery, radiation therapy and clinical outcome [PGD*06] [GWF*05]. Our framework aims to facilitate this assessment before and during treatment.

For many patients, a set of two MRI scans (T1- and T2-weighted) and one PET scan is produced as part of the diagnosis. The resulting datasets are often evaluated by overlaying 2D slices, or by comparing 2D slices side by side. However, in recent years the segmentation of three-dimensional regions of interest such as tumor tissue has been greatly simplified by the introduction of automatic segmentation techniques such as the approach presented by Corso et al. [CSD*08]. By separating voxels within a 3D dataset having a high probability of belonging to a tumor region from voxels representing healthy tissue, a binary mask dataset is generated which has non-zero values only where the image indicates tumor tissue. We call these regions *segmented tumor regions* (STRs). The system has been designed with the objective to subsequently facilitate the understanding of the spatial extent of the generated three-dimensional STRs in relation to each other. By exploiting comparative visualization strategies, the three different standard imaging modal-

DOI: 10.2312/PE.VMV.VMV13.105-112

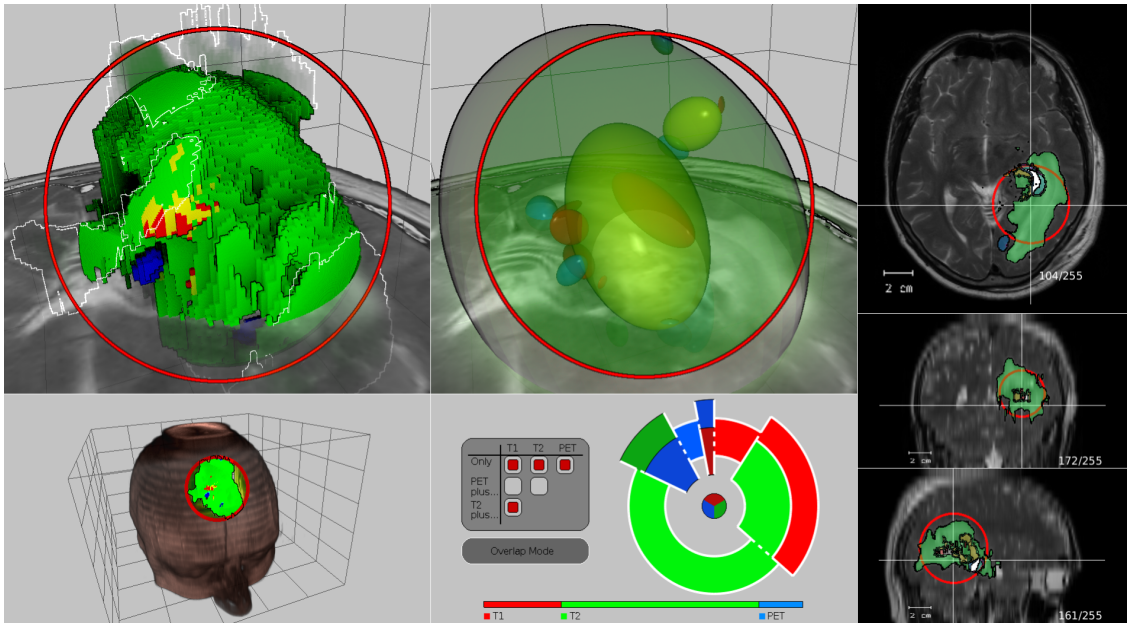

Figure 1: *An overview of our framework. In the upper left, the tumor intersection closeup with the context view below. In the upper middle, the tumor ellipsoid diagram; in the lower middle, the volume overlap diagram. At the right, three linked 2D slice views showing the STRs and their overlaps, as well as the respective 2D projection of the tumor lense.*

ities can be compared to each other simultaneously to allow a clinical characterization of the tumor. Areas where two or even three STRs intersect can be easily identified, which may indicate the center of the actual tumor tissue.

Besides an initial diagnosis, an optimal application is required to assess the progress of the tumor over time to track development of the tumor disease, discover recidivisms, make treatment decisions and carry out therapy such as surgery or radiotherapy. Therefore, our framework also aims to visualize the extension or the shrinking of different tumor parts between different stages of treatment to determine potential recidivisms and tumor parts that are inactive. Surrounding anatomical structures are visualized to determine the direction in which the tumor is proliferating.

2. Related work

In recent years, much work has been published concerning multimodal volume rendering of brain imaging data. However, most publications focus on pre-operation planning, using context modalities such as fMRI to find optimal access paths instead of comparing the size and shape of multimodal segmentations of the same tumor for diagnosis and treatment planning. An early example of multimodal brain image rendering in 3D is given by Serra et al. [SKG*98], who simultaneously integrate CT, MRA and MRI modalities into a virtual workbench using a slicing approach. Beyer et al. [BHWB07] introduce a framework that supports planning of the optimal incision point by using a skull-peeling algorithm, as well as multimodal visualization of the underlying structures. Herghelegiu et al. [HMP*12] present a system for optimal planning of a brain tissue biopsy which helps to avoid hitting blood vessels or other critical areas. Jannin et al. [JFS*00] extract silhouettes of multimodal data to blend over a context volume to visualize multiple structures, which inspired the use of silhouettes and contours in this paper as a 3D shading approach.

Rieder et al. [RRRP08] presented interaction methods to aid surgeons in simultaneously perceiving multimodal data to support intervention planning. Rieder et al. [RSHP08] introduced a framework for blending multimodal MRI images which used specialized clipping and illumination to separate tumor tissue from surrounding brain tissue. Furthermore, Rieder et al. [RWS*10] introduced the tumor map to visualize the treatment success of liver tumors with percutaneous radiofrequency ablation. This map provides a pseudo-cylindrical representation of the tumor surface, which simplifies the visualization of the tumor structure.

Tietjen et al. [TMS*06] presented techniques to show the extent of multiple segmented objects within 2D slices by attaching vertical bars representing the length of each object next to a 2D slice of a context volume. Viola et al. [VKG04] introduced a technique which suppresses less relevant parts of the data by mapping data importance to visibility. A similar approach is used for the tumor lense in our paper.

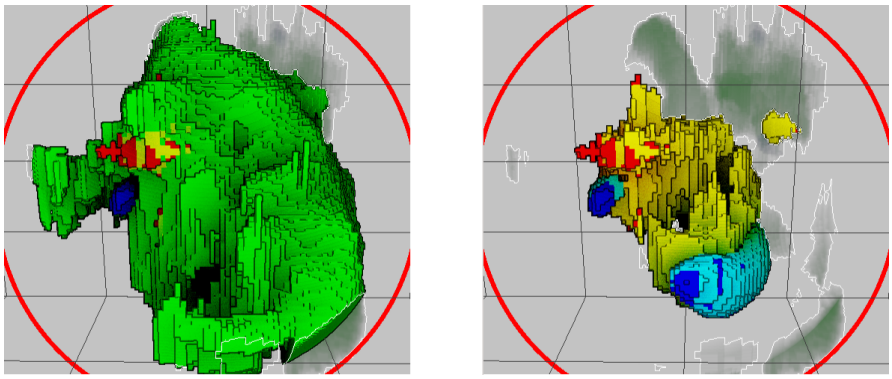

(a) (b)

Figure 2: *The tumor intersection closeup. The region where only the T2-weighted MRI modality is active (green) occludes other segments (a). By deselecting it (b), the underlying STR configuration can be made visible. In this case, relatively large areas are displayed where the T2-weighted MRI and the T1-weighted MRI modality (yellow) as well as the T2-weighted MR and the PET modality (turquoise) are active. Note that contours and ambient occlusion convey the shape of the regions well.*

3. 3D visualization

As large sections of the different tumor modalities will overlap and therefore occlude each other, the challenge lies in providing the user with a spatially comprehensible image. Our central 3D visualization of the tumor uses a modified GPU ray casting approach, where rays traverse a volume dataset in which the co-registered STR masks are bitwise encoded as integer data, the so called *tumor index*. Thus, for instance, a single 8 bit dataset is able to hold information about three different tumor masks at two different points in time. The color scheme used during raycasting when hitting an STR is based on the RGB color space, where each of the three STRs from the standard image modalities is represented by the respective color channel. The T1-weighted MRI STR is associated with the red channel, the T2-weighted MRI STR with the green channel and the STR segmented from the PET modality with the blue channel. For overlaps between the STRs, an area of intersection of the first and the second STR would be assigned the color yellow, whereas a region where only the third mask was active would be assigned the color blue. This basic color scheme was chosen instead of conventional transfer functions as we want to visualize a set of binary masks and their intersections, rather than a distribution of intensity values.

3.1. Surface shading model

To perceive the surface, shape and size of objects, shading is a valuable tool in volume rendering [LR11], especially when having to distinguish overlapping objects. However, depending on the segmentation method used as well as on the image modality, the different STRs are often very irregular and frayed, making an intensity gradient-based shading approach difficult to apply. Instead, we chose a volume illustration approach to shading, in which depth-based edge detection is used to generate and display silhouettes and contours of each STR. First, a raycasting image is generated using nearest instead of linear or even more complex filtering methods. We chose this simple filtering model as we try to visualize the segmentation masks themselves and not a reference volume with the help of segmentation masks. As a result, the rendering resembles the original STR data more closely. On the rendered tumor surface, neighboring voxels whose distance from the viewer differ even only by the extent of a single voxel will have depth values which differ by a small, but significant amount due to the simple, coarse filtering. A subsequent depth-based edge detection on the raycasting image using depth-based line thickness brings out these small differences as subtle contours, resulting in a ”building brick effect” (see Fig. 2). Silhouettes are emphasized with black lines. Inner contours are kept more subtle, using very thin lines in a light gray color. Additionally, screen space ambient occlusion as suggested by Shanmugam and Arikan [SA07] is used to further illustrate differences in relative depth by displaying shadow-like halos around STR areas in the foreground. Our implementation of this technique relies only on depth values and is therefore suitable for the non-smooth surface structure of the STRs. The result of the combination of ambient occlusion with contour and silhouette lines is an unobtrusive, yet effective shading model conveying the size, shape and relative depth of visible regions and surfaces when applied to the coarse STR surfaces (see Fig. 2).

3.2. Tumor context view

The lower of the two 3D views (see Fig. 1, left column) displays a context volume together with the tumor data, generated using ray casting with a user-defined transfer function

for the context volume. We call this view the *tumor context view* (TCV). In order to always be able to see the tumor region, the ray casting algorithm has been modified for this view. Before rendering, the axis aligned bounding box of the combined STRs is calculated and saved in a separate proxy geometry. If a ray is going to hit this tumor box, the region which the ray traverses within the context volume before it hits the box is discarded. Within the tumor box, only tumor data is rendered with full opacity while the opacity of the context volume is set to 0. If the ray has not hit tumor tissue on its way through the tumor box, the ray continues through the context volume, now with the regular opacity given by the transfer function. Thus, the tumor box is always visible showing the entire union of all STRs, independently of the rotation of the context volume (see Fig. 1, bottom left).

3.3. Tumor intersection closeup in single time step mode

In single time step mode, the upper 3D view displays the combined STRs of the three standard image modalities recorded at the same time in the so-called *tumor intersection closeup* (TIC). For this view, we use a two level ray casting process [HBH03] which has been modified to deal with the issue of overlapping STRs. We distinguish two regions: an inner spherical region of interest which we call the *tumor lense*, initially located around the center of mass of the union of all STRs, and the outer context region. As a ray traverses the volume, we check whether or not the current sample point lies within the spherical tumor lense. If not, a low opacity is used to visually suppress regions that are currently not of interest. If the ray sample lies within the tumor lense and tumor tissue is hit, all previously composited color from the outer region is discarded to provide an unobscured view of the lense content. However, the depth of the first hit point of the outer region is always saved separately to provide contour lines of context STRs which would currently occlude the tumor lense. These outer contour lines are overlaid in the final image as a subtle hint that there are currently occluded areas (see Fig. 2). Then, the local tumor index is looked up. The user can control which STRs should be visible within the spherical region by enabling a set of modalities or modality intersections in the tumor overlap diagram (see Section 5). If the retrieved tumor index is part of the current user selection, a high opacity is used, otherwise the found STR voxel remains invisible.

In order to better navigate this view, the camera view vector is synchronized with the one used in the TCV. When rotating the tumor in TIC view, the camera orientation (but not position) is synchronized to the TCV camera. This enables the user to track from which direction relative to the patient's head the tumor closeup is seen. For a further visual connection, a red ring is drawn around the region of interest representing the radius the user has currently selected for the tumor lense in both TCV and TIC (see Fig. 1). This ring can be used to control the radius of the lense by grabbing and dragging it with the mouse cursor. The center of the tumor lense can be changed by the user by dragging the mouse within the 3D view while pressing an additional modifying key. This also changes the trackball center and camera focus of the TIC which always follows the center of the tumor lense, keeping the navigation unambiguous.

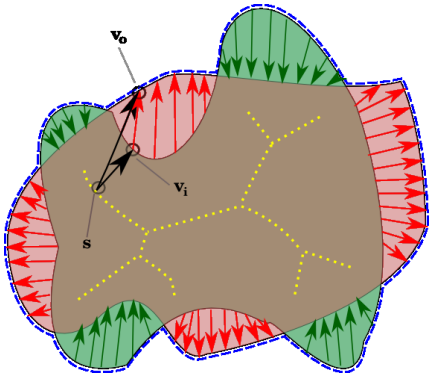

Figure 3: *The image illustrates the method to find the changed surface of an STR between two points in time. For each voxel v_o on the outer surface of the union of the two STRs (blue), the closest of all skeleton points (yellow), s, is determined. Then, we consider the set of all voxels on the surface of the respective other STR to which s is the closest skeleton point. We search this set for the voxel v_i which is closest to v_o (see black arrows). The vector connecting the two represents the change of the tumor surface.*

3.4. Tumor surface closeup in consecutive time steps

Instead of focusing on different STRs recorded within the same time frame, the user may also select the same STR recorded at different points of time t_1 and t_2. The so-called *tumor surface closeup* can be used to analyze regions in which the tumor has changed due to treatment and to monitor how each imaging method displays these changes. This visualization uses a pre-process to identify regions in which the selected STR has shrunk, grown or remains unchanged. During this process, the distance between the surface voxels of STR_{t_1} and STR_{t_2} is determined.

First, the intersection $STR_{\cap}$ of STR_{t_1} and STR_{t_2} is calculated. We then use the technique introduced by Cornea et al. [CSYB05] to compute a set of points STR_{skel} within $STR_{\cap}$, representing its inner curve-skeleton. The skeleton voxels will then be used to determine the direction of change from each surface voxel of STR_{t_1} and STR_{t_2}. Since the STRs can have complex elongated forms including concavities, computing the change in the direction of the center of mass, for instance, may not be sufficient. Furthermore, due to the fuzziness of the MRI data, we did not want to rely directly on intensity gradients to calculate the direction of surface change of the STRs. The skeleton provides inner reference

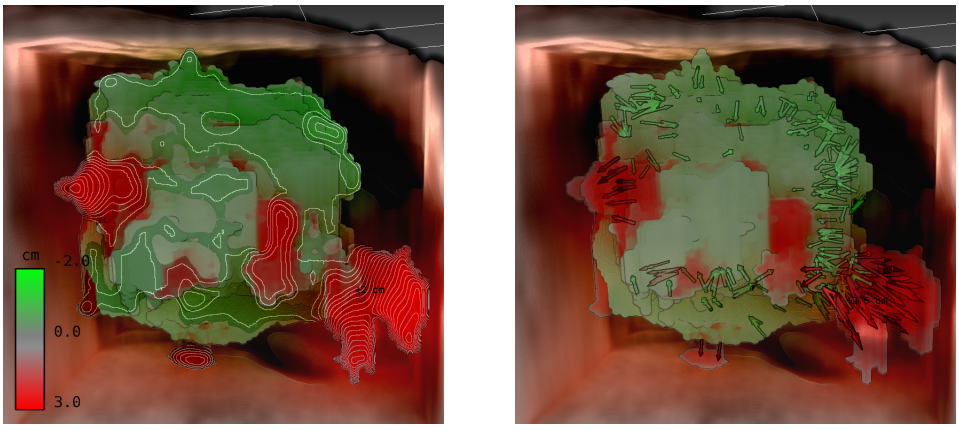

(a) (b)

Figure 4: *The tumor surface closeup when comparing consecutive time steps. After skeletonization and the distance calculation, the user can view the resulting changes on a height map (a) which is projected onto the outer union of the two STRs. In this case, the STR has mainly shrunk (green areas), but has also developed two extensions which were not there before (red areas). As an alternative to the height field, the user can view the surface changes between time steps using arrows (b).*

points without using gradients, while observing the potentially complex shape of the STR.

Next, we find the set of outer surface voxels STR_o of the union of STR_{t_1} and STR_{t_2}, i.e., all voxels belonging to a surface of either STR_{t_1} or STR_{t_2} where the respective other STR is inactive (see Fig. 3). For each point in STR_{skel}, we separately calculate the closest surface voxel belonging to STR_{t_1} and STR_{t_2}. Conversely, for each voxel in STR_o, we calculate the closest skeleton point in STR_{skel}. Afterwards, for each voxel v_o in STR_o, we look up the corresponding closest skeleton point and search the associated set of closest surface voxels from the respective other surface for the closest voxel v_i to v_o. Thus, we have found a vector between the two surfaces of the chronologically separated STRs representing the direction and amount of change towards or away from the skeleton of their intersection. An additional flag is saved to indicate whether v_o originally belonged to STR_{t_1} or STR_{t_2}, representing tumor recession or tumor growth, respectively.

In order to visualize the previously calculated surface changes, the user of our framework can choose from two visualization methods. Both use ray casting as a basis, with additional overlaid information. In the first method, for each image pixel, the ray's first hit point is used to look up the associated directional change information. The length of the surface change vector and the type of change (growth or shrinkage of the tumor) is then saved to a separate image target. After ray casting, a Gauss filter is applied to that target in order to smooth the resulting image. The smoothed image is then transformed to a height field image, with the color ranging from transparent to opaque green for regions of tumor shrinkage and from transparent to opaque red for regions of tumor growth. Regions where the STR has not changed remain transparent. The height field image is then blended over the original ray casting image (see Fig. 4(a)). By moving the mouse cursor over the resulting image, the unfiltered value of the length of the surface change vector is looked up from the ray casting result. The corresponding distance is displayed next to the cursor to reveal the exact local temporal difference. Instead of the height field, the user can also choose a geometrical representation of tumor surface change which uses green or red arrows (see Fig. 4(b)). The arrows are generated on the GPU using geometry shaders to allow for interactive frame rates. The number of shown arrows can be limited using length thresholding to avoid cluttering and to focus on areas of large surface changes.

Finally, the resulting STR surface rendering is embedded in a visualization of one of the datasets which were used to generate the STR. Using the same approach as in the tumor context view (see Section 3.2), this dataset containing surrounding anatomical structures of interest can be rendered simultaneously by exploiting a ray casting approach which differs between a context volume and the STRs (see Fig. 4(b)). This allows the user to track directly in which direction relative to the brain the tumor has extended or receded, and whether a possible tumor extension is close to sensitive brain regions which may require additional treatment.

4. 2D visualization

In order to support the 3D visualization methods, our framework includes traditional 2D slice visualizations of each anatomical plane, showing linked transversal, coronal and sagittal slices of the context volume. In each 2D view, the STRs are overlaid using the same color scheme as in the

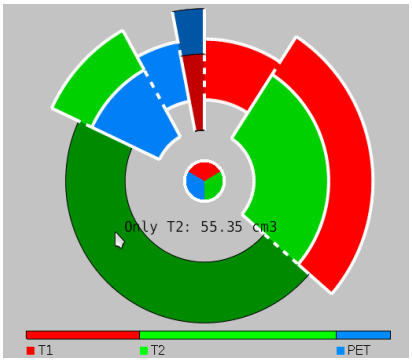

Figure 5: *The tumor overlap diagram shows overlap sizes within the STRs as well as the original size of any STR compared to the others. Angle sizes represent the size of the overlaps, while areas of the same color represent the size of an STR without considering overlaps. Overlaps currently visible within the tumor lense are bordered in white. Stippled lines are used to indicate borders between segments which are part of the same STR.*

3D visualizations (see Fig. 1, right column). Additionally, the radius and center of the tumor lense (see Section 3.3) are projected onto each 2D view and displayed as a circle in single time step mode, similar to the one used in the corresponding TIC and TCV. By dragging the mouse over the TIC, the 2D views are updated to match the current 3D location of the cursor. This allows an interactive exploration of the exact outer structure of the overlapping STRs in 3D, while relying on the more unambiguous 2D views. To further visually link the 2D and 3D views, the transversal slice currently selected by the user can optionally be included in the TIC as well as the tumor overlap diagram (see Section 5) as a semi-transparent plane. This allows the user to connect anatomical structures visible in the slice to the more complex 3D representation of tumor tissue in those visualizations (see Fig. 1, upper left / middle).

5. Tumor overlap diagram

The *tumor overlap diagram* depicts the relative sizes of the three currently displayed STRs (see Fig. 5). It distinguishes between three kinds of regions:

- Only one modality is active (non-stacked ring segments)
- Two modalities are active (stacked ring segments)
- All modalities are active (circle in middle of diagram)

The diagram was designed to enable the user to compare the relative sizes of the STR overlaps using the opening angle of the segments. At the same time, the relative sizes of the STRs themselves are preserved by the color distribution within the diagram. For instance, by comparing the area of diagram regions colored in green with the other two colors in Fig. 5, the user can quickly determine that the T2 STR is roughly twice as large as the T1 STR and five times as large as the PET STR. However, judging by the radian measure of the non-stacked green ring segment, it is also obvious that a large area of the T2 STR is active where none of the other STRs show tumor tissue. Based on this observation, the user may want to mask this region in the 3D views and concentrate on those areas where two or more STRs coincide.

For this reason, the tumor overlap diagram is the central tool in our framework to select the currently active STRs and their overlaps. By selecting and deselecting segments of the diagram with the mouse, the user can specify an arbitrary configuration of active segments. When viewing a single time step, this will cause the tumor lense in the TIC to only show the active segments. For instance, this enables the user to only show the three-way overlap section along with the section where only the PET STR is active. This allows for a flexible comparison of all STR combinations. Additionally, when moving the mouse over each segment, the size of the corresponding segment is displayed to provide a quantitative assessment of the STRs.

6. Tumor ellipsoid diagram

As mentioned before, the structure of the STRs and the resulting overlaps is mostly irregular and not easily comprehensible when rendered in 3D due to the challenge of depth perception and occluded structures. While looking for a simplified representation of the STRs in 3D, we found that despite the irregularity of the tumor areas, their shape is generally still convex enough to be approximated by a closest fitting ellipsoid. Therefore, the *tumor ellipsoid diagram* provides a simplified view of the tumor area by replacing the original, arbitrarily shaped STRs with ellipsoids while following the same color scheme (see Fig. 6(a) and (b)). The result is similar to a 3D Venn diagram, which facilitates a quick understanding of the rough structure of the STRs. The point of view from which the diagram is seen is synchronized to the TIC, as well as the currently visible STRs or STR overlaps as specified by the user in the tumor overlay diagram. Thus, the tumor ellipsoid diagram can be used as a point of reference to better comprehend the more detailed rendering provided by the TIC.

To compute the ellipses, a pre-processing step is necessary. The voxels belonging to each STR as well as all occurring overlapping areas between the individual STRs are first divided into their connected components. This division is necessary since a single ellipsoid is not sufficient to approximate STRs which are divided into several separate active regions, which occurs often within STRs derived from PET modalities. We then use the algorithm proposed by Khachiyan [Kha96] to compute the closest fitting ellipsoid for each set of voxels within a connected component. After this process is complete, the resulting ellipsoid centers, radii and rotation matrices are stored together with the STR dataset.

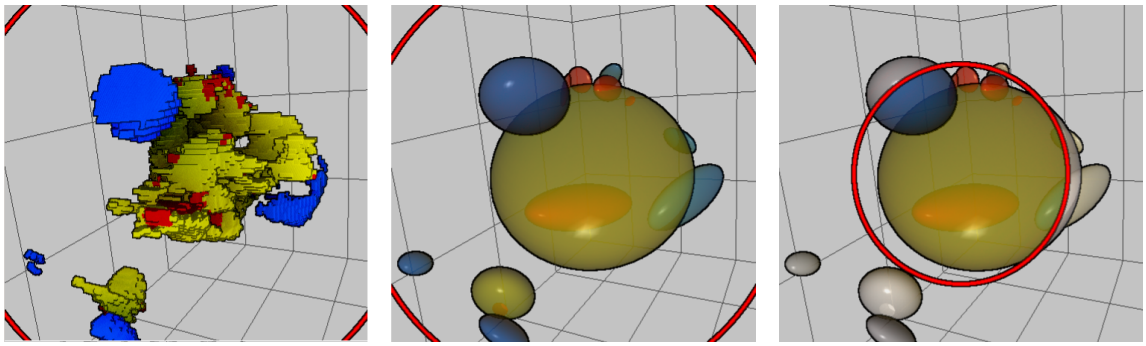

(a) (b) (c)

Figure 6: *The tumor ellipsoid diagram uses the closest fitting ellipsoids for each connected component of each currently active STR (a) to provide a simplified view of the tumor data (b). This diagram facilitates the spatial understanding of the STRs, as it avoids the visual cluttering and occlusions which may occur when applying ray casting on the original data. The tumor lense radius is visualized by greying out ellipsoid regions which are outside of the lense (c).*

To generate the diagram, the ellipsoids are rendered as geometric primitives, without the use of raycasting. The user can control the transparency of each visible ellipsoid. To guarantee correct order independent transparency, we employ an A-buffer as introduced by Carpenter [Car84] to store a list of fragments per pixel which is used after rendering to correctly sort the stored fragments based on depth and correctly blend individual ellipsoids over each other. The screen space ambient occlusion shading used in the TIC is also employed for this diagram, as it enables the user to better judge the size and relative depth of the different ellipsoids. Additionally, the ellipsoid surfaces are shaded using the Phong reflection model with the light source located at the point of view of the user to convey their 3D shape. The red ring which indicates the current radius of the tumor lense is also displayed in this diagram to support the visual connection between the views. Regions of ellipsoids which are outside of the lense are displayed in shades of grey to distinguish them from regions within the lense (see Fig. 6(c)).

7. Domain expert feedback

We demonstrated our framework to two domain experts familiar with the process of diagnosis and treatment of brain tumors. Both agreed that our framework could be an overall useful visual tool. One complimented the use of nearest filtering in the 3D raycasting views and the subsequent ability to exactly identify regions in the TIC with only few voxels of tumor tissue which might nevertheless require treatment.

Both experts agreed that the tumor overlap diagram required some adaptation to its approach of presenting information with circular elements as opposed to more traditional bar charts. However, after a short learning phase, one domain expert was able to judge within seconds which modalities and which overlaps were dominant in a given dataset. One of the experts suggested the implementation of an additional interactive table to control the visibility of tumor overlaps, as it shows more clearly what the individual ring segments signify exactly (see Fig. 1, middle bottom). Another suggestion which we adopted was to include an optional non-overlap mode in the diagram which allows the user to select STRs without considering overlaps, which results in simplified visualizations in the TIC as well as the tumor ellipsoid diagram and provides the user with a less complex overview of the original data. Furthermore, a rectangular bar that displays the STR sizes without overlaps next to each other was added below the tumor overlap diagram in order to provide a more conventional point of reference (see Fig. 5). The tumor ellipsoid diagram was described by both as useful for explaining a diagnosis to a patient rather than the diagnosis itself, since the convexity of the ellipsoids causes tumor overlaps to be exaggerated at times. As one of the experts noted that the shape and relative position of the ellipsoids could be indistinguishable depending on the camera position, we added a simple form of shading to the diagram (see Section 6).

Finally, the tumor surface closeup in consecutive time steps was commended by both experts as being useful to judge the change of a tumor over time. This task is challenging due to the potentially complex structure of brain tumors as well as the inhomogeneity of the changes a tumor may undergo. The experts welcomed our approach as it shows complex information of the whole tumor in one simple to understand visualization.

8. Conclusion and future work

We have presented an interactive framework for the exploration of tumor segmentations derived from multimodal brain images. Our framework provides the user with an exact understanding of the inner structure of the union of all STRs.

The user can specify arbitrary subsets of overlaps or areas where only one STR is active. These subsets are directly visible since areas that are currently not of interest for the viewer are removed. Our shading approach supports a quantitative understanding and comparison of the shape and size of different kinds of imaging modalities. Exact size and distance values are available to the user in different views due to information overlays displayed next to the mouse cursor. To provide an overview of the STR structure, the user can refer to the tumor ellipsoid diagram which reduces spatial complexity by providing an approximate 3D Venn diagram. The tumor overlap diagram provides an instant overview of the entire size of each STR as well as the size of overlaps with other segmented regions. As the area of a color in the diagram directly translates to the size of the associated STR, it could be a powerful tool when used as a kind of glyph to view many datasets next to each other, for example to visually determine relations between the size of the STR and the imaging method.

We demonstrated our framework to domain experts and used their feedback to improve our visualizations. By modifying our color scheme as well as extending the tumor overlap diagram, it may be possible to visualize more than three modalities at the same time. Our system could also be extended to include other kinds of modalities such as fMRI and be subsequently used for surgery planning.

Acknowledgements

This work was partly supported by grants from the Deutsche Forschungsgemeinschaft (DFG), SFB 656 MoBil Münster, Germany. The presented concepts have been integrated into the Voreen volume rendering engine (*voreen.uni-muenster.de*).

References

[BHWB07] BEYER J., HADWIGER M., WOLFSBERGER S., BUHLER K.: High-quality multimodal volume rendering for preoperative planning of neurosurgical interventions. *IEEE Transactions on Visualization and Computer Graphics 13* (2007), 1696–1703. 2

[Car84] CARPENTER L.: The A-buffer, an antialiased hidden surface method. *SIGGRAPH Comput. Graph. 18*, 3 (Jan. 1984), 103–108. 7

[CSD*08] CORSO J., SHARON E., DUBE S., EL-SADEN S., SINHA U., YUILLE A.: Efficient multilevel brain tumor segmentation with integrated bayesian model classification. *Medical Imaging, IEEE Transactions on 27*, 5 (may 2008), 629 –640. 1

[CSYB05] CORNEA N. D., SILVER D., YUAN X., BALASUBRAMANIAN R.: Computing hierarchical curve-skeletons of 3d objects. *The Visual Computer 21* (2005). 4

[GWF*05] GROSU A. L., WEBER W. A., FRANZ M., STÄRK S., PIERT M., THAMM R., GUMPRECHT H., SCHWAIGER M., MOLLS M., NIEDER C.: Reirradiation of recurrent high-grade gliomas using amino acid pet (spect)/ct/mri image fusion to determine gross tumor volume for stereotactic fractionated radiotherapy. *Int J Radiat Oncol Biol Phys 63*, 2 (2005), 511–9. 1

[HBH03] HADWIGER M., BERGER C., HAUSER H.: High-quality two-level volume rendering of segmented data sets on consumer graphics hardware. In *Proceedings of the 14th IEEE Visualization 2003 (VIS'03)* (Washington, DC, USA, 2003), VIS '03, IEEE Computer Society, pp. 40–47. 4

[HMP*12] HERGHELEGIU P., MANTA V., PERIN R., BRUCKNER S., GRÖLLER M. E.: Biopsy planner - visual analysis for needle pathway planning in deep seated brain tumor biopsy. *Computer Graphics Forum 31*, 3 (June 2012), 1085–1094. 2

[JFS*00] JANNIN P., FLEIG O. J., SEIGNEURET E., MOR X., RAIMBAULT M., CEDEX R.: Multimodal and multi-informational neuronavigation. In *In Proc. of CARS - Computer Assisted Radiology and Surgery* (2000), pp. 167–172. 2

[Kha96] KHACHIYAN L. G.: Rounding of polytopes in the real number model of computation. *Math. Oper. Res. 21*, 2 (May 1996), 307–320. 6

[LR11] LINDEMANN F., ROPINSKI T.: About the influence of illumination models on image comprehension in direct volume rendering. *IEEE TVCG (Vis Proceedings) 17*, 12 (2011), 1922–1931. 3

[PGD*06] PIROTTE B., GOLDMAN S., DEWITTE O., MASSAGER N., WIKLER D., LEFRANC F., BEN TAIB N. O., RORIVE S., DAVID P., BROTCHI J., LEVIVIER M.: Integrated positron emission tomography and magnetic resonance imaging-guided resection of brain tumors: a report of 103 consecutive procedures. *Journal of Neurosurgery 104*, 2 (2006), 238–53. 1

[RRRP08] RIEDER C., RITTER F., RASPE M., PEITGEN H.-O.: Interactive Visualization of Multimodal Volume Data for Neurosurgical Tumor Treatment. *Computer Graphics Forum (Special Issue on Eurographics Symposium on Visualization) 27*, 3 (2008), 1055–1062. 2

[RSHP08] RIEDER C., SCHWIER M., HAHN H. K., PEITGEN H.-O.: High-quality multimodal volume visualization of intracerebralral pathological tissue. In *Proc. VCBM* (2008), pp. 167–176. 2

[RWS*10] RIEDER C., WEIHUSEN A., SCHUMANN C., ZIDOWITZ S., PEITGEN H.-O.: Visual Support for Interactive Post-Interventional Assessment of Radiofrequency Ablation Therapy. *Computer Graphics Forum (Special Issue on Eurographics Symposium on Visualization) 29*, 3 (2010), 1093–1102. 2

[SA07] SHANMUGAM P., ARIKAN O.: Hardware accelerated ambient occlusion techniques on GPUs. In *Proceedings of the 2007 symposium on Interactive 3D graphics and games* (New York, NY, USA, 2007), I3D '07, ACM, pp. 73–80. 3

[SKG*98] SERRA L., KOCKRO R. A., GUAN C. G., LEE N. H. E. C., LEE E. C. K., LEE Y. H., CHAN C., NOWINSKI W. L.: Multimodal volume-based tumor neurosurgery planning in the virtual workbench. In *in the Virtual Workbench, Proc. IEEE VRAIS'98* (1998), pp. 167–173. 2

[SPS*10] SHERMAN J., PREVEDELLO D., SHAH L., RAGHAVAN P., POURATIAN N., STARKE R., LOPES M., SHAFFREY M., SCHIFF D.: Mr imaging characteristics of oligodendroglial tumors with assessment of 1p/19q deletion status. *Acta Neurochirurgica 152* (2010), 1827–1834. 1

[TMS*06] TIETJEN C., MEYER B., SCHLECHTWEG S., PREIM B., HERTEL I., STRAUSS G.: Enhancing Slice-based Visualizations of Medical Volume Data. In *Proceedings of the Eurographics / IEEE VGTC Symposium on Visualization* (2006), Fellner D., Möller T., Spencer S., (Eds.), EGPub, pp. 123–130. 2

[VKG04] VIOLA I., KANITSAR A., GRÖLLER M. E.: Importance-driven volume rendering. In *Proceedings of IEEE Visualization'04* (2004), pp. 139–145. 2