Eurographics Workshop on Visual Computing for Biology and Medicine (2010)
D. Bartz, C. P. Botha, J. Hornegger, R. Machiraju, A. Wiebel, and B. Preim (Editors)

# Estimation of the Midsagittal Plane for Sideness Determination of Malignant Structures of Head and Neck

Ivo Rössling[1] and Peter Hahn[2] and Christian Cyrus[1] and Lars Dornheim[2]

[1]Otto-von-Guericke University, Magdeburg, Germany
[2]Dornheim Medical Images, Magdeburg, Germany

**Abstract**

*Besides other main criterias like size, infiltration and anatomical district, the sideness of tumor and local metastases (i. e., malignant lymph nodes) is very crucial for head and neck tumor assessment. An automatic sideness determination may speed up tumor staging noticeably. Inspired by preliminary work of others we present a modified approach for the estimation of the midsagittal plane based on surface meshes. The suitability of the computed result being used in the course of TNM classification was tested in a case study, in terms of an according automatic determination of the sideness for segmented lymph nodes and tumors.*

Categories and Subject Descriptors (according to ACM CCS): I.4.8 [Image Processing and Computer Vision]: Scene Analysis ▷ Shape

## 1. Introduction

The assessment of anatomical and pathological structures based on imaging techniques plays a vital role in surgical diagnostics and therapy planning. Looking for potential abnormalities, medical doctors exert their through the years acquired expert knowledge to detect discrepancies between a given acquisition of a region of interest and its typical appearance within the respective modality. Such aberrances in particular become manifest in an alteration of coloration, homogeneity or shape. But also atypical dissymmetries may indicate potential anomalies.

The median or midsagittal plane constitutes a special reference in this regard, in particular in respect of growing computer-assisted diagnostics. It can, for example, be used for alignment with respect to some reference system (like the Talairach frame [TT88], for instance) or provide a basis for registration (either of similar structures or against some given atlas). It may also serve as starting point for a dissymmetry analysis, like it is done in Mancas et al. [MGM05], for an automatic localization of potential tumor areals in the head region.

In the context of segmented anatomical and pathological structures, a reasonably good estimation on location and orientation of this plane in many cases already allows a correct determination of the sideness for several individual structures in a fully automatic manner. This is in particular important in the course of tumor staging (TNM classification), since the laterality of overall lymph node affection is crucial for finding and assessing suitable treatment and therapy options (cf. Figure 1). Laterality here refers to the sideness of affected lymph nodes relative to the tumor: ipsilateral (same side), contralateral (opposite sides), unilateral (only one side, either ipsi- or contralateral), bilateral (both sides). This particular use case shall be the main focus of the following study.

## 2. Related Work

The issue of symmetry determination has been considered repeatedly in the pertinent literature on general and medical image processing. In the medical context the main focus of interest has thereby been on the brain. A good survey on this topic is provided by Prima et al. **[POA02]**.

For the determination of such a symmetry plane, two basic strategies can be distinguished which arise from the simultaneous anatomical and symmetrical nature of the data. On the one hand, a search for specific landmarks or other morphological features with subsequent fitting can be conducted. Brummer [Bru91], for instance, applies the Hough transform principle to detect the interhemispheric fissure.



DOI: 10.2312/VCBM/VCBM10/109-115

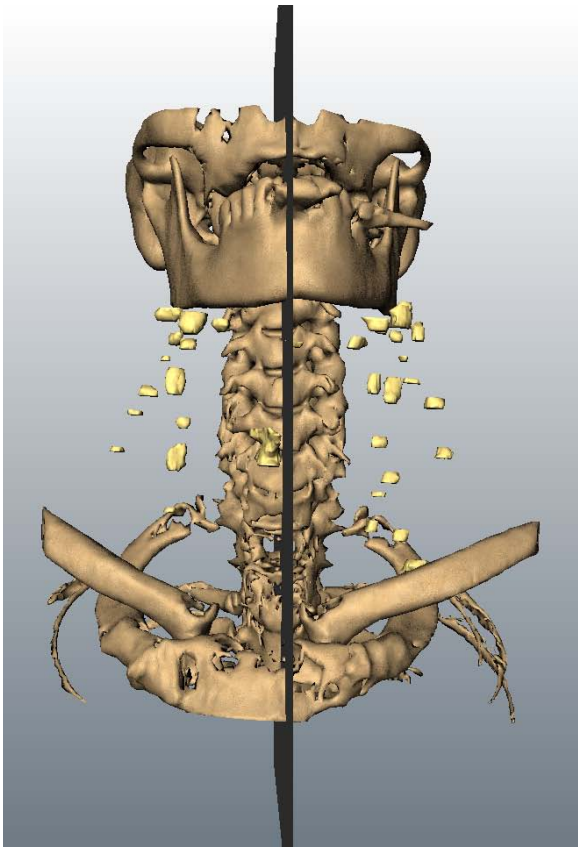

**Figure 1:** *Result for the computation of an estimate for the midsagittal plane, based on a given bones' surface mesh. If the estimate is reasonably good, it allows for automatic determination of the sidnesses of other anatomical structures, such as tumor (middle) and lymph nodes (left and right).*

On the other hand, the objective can be the determination of a plane that (more or less) maximizes some respective symmetry criteria. Common approaches in this regard are non-rigid registration of a (3*d*) dataset against ists mirror image (cf. Prima et al. [POA02]) or the computation of a matching by maximizing the cross-correlation under rotation and translation (cf. Liu et al. [LCR01]). Also histograms were sometimes used as pivotal criteria (cf. Mancas et al. [MGM05]). In all these cases gray-level or edge images were used as data.

Alternative methods range from using statistical measures (cf. Thirion et al. [T*98]) over algebraic approaches (cf. Keller et al. [KS04]) up to so-called *generalized symmetry transforms* (cf. Reisfeld et al. [RWY95]) and *reflective symmetry descriptors* (cf. Kazhdan et al. [KCD*03]).

The idea of using symmetries in the orientation histogram for identifying symmetries in the original (2*d*) image space was first presented by Sun [Sun95]. In the optimal case, the reflective symmetry axis of the original image is also a symmetry axis in the (circular) orientation histogram along which the two half histograms should match each other by reflection. A least-squares fitting of the histogram against its stepwise rotated mirror image should therefore give a good approximation of the symmetry axis' orientation.

This basic idea can not only be extended from 2*d* to 3*d*, it can also easily be adapted to arbitrary surface meshes instead of only gray-level image data. Sun et al. [SS97] discuss these use cases exemplarily. However, they do not examine quality nor reflect on usability of the obtained result (in terms of the symmetry plane) subsequently.

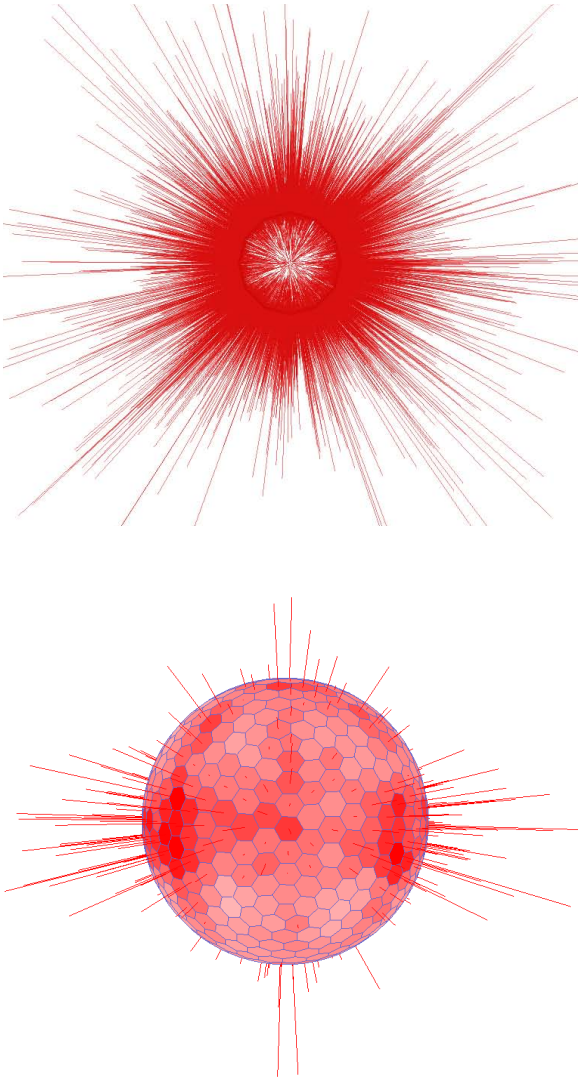

**Figure 2:** Top: 3*d EGI for bones' surface mesh shown in Figure 1*. Bottom: *Corresponding* 3*d orientation histogram. To account for the viewing angle, the bin value is double-coded, in the length of the peak and in the color of the face.*

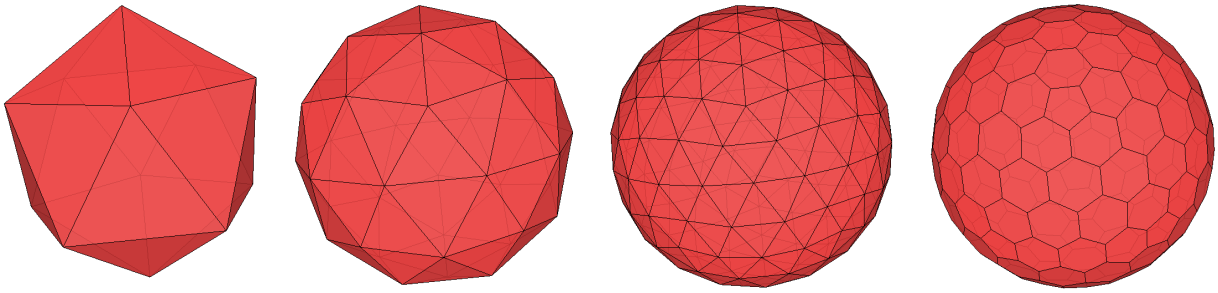

**Figure 3:** *Tessellation of the unit sphere to yield the bins for a* 3*d spherical histogram.* Left*: Initial icosahedron.* Middle*: After first and second subdivision.* Right*: After joining triangular faces' centerpoints to hexagons and pentagons.*

## 3. Materials and Methods

For our investigations we used segmentations of neck and head in terms of surface triangle meshes. Voxel-based segmentations in this regard can easily be transformed into equivalent triangle meshes without significant extra cost (using e. g. Marching Cubes or related methods). The goal was to compute an estimation for the midsagittal plane based on a given segmentation of the bones of head and neck. For this resulting plane, the correctness of an automatic sideness determination of the remaining segmented structures should be evaluated.

Regarding the computation of midsagittal plane we basically oriented ourselves on the approach due to Sun et al. [SS97]. First of all, a principal component analysis is performed to compute barycenter and main axes of the segmented bones.

This step is motivated by the fact that in the (theoretical) case of perfect symmetry it holds (cf. Minovic et al. [MIK92]):

- *Any plane of symmetry of a body is perpendicular to one of the principle axes.*
- *Any axis of symmetry of a body is a principle axis.*

In the first instance, however, it is not clear which of the principle axes is the respectively qualified one for our aim.

To determine the orientation of the symmetry plane, Sun et al. [SS97] in the next step map the surface mesh's face normals onto the unit sphere. To account for the fact that the mesh may not necessarily be homogeneously meshed, each normal is weighted by the area of the corresponding face. One obtains the so-called *Extended Gaussian Image* (EGI, see Figure 3, top), which is subsequently being discretized to a spherical 3*d* orientation histogram (Figure 3, bottom).

This kind of discretization is done by tessellating the unit sphere with some convex polytope of desired resolution. Any face or cell of the polytope corresponds to one bin of the histogram. Each vector of the EGI has a direction and a length. The direction defines a ray that intersects the polytope in one of its faces and thus determines the target histogram bin. The length of the vector determines the value that is contributed to that bin. Regarding the tessellation of the sphere it is desirable to achieve an (almost) regular polytope. Unfortunately, perfect regularity is not possible, except for the five platonic solids. A good approximation, however, can be achieved by recursively subdividing the triangular faces of an icosahedron (see Figure 3). In each recursion step, any triangle is split into four smaller triangles by joining the midpoints of its three sides with a new edge for each pair. Then the new vertices are projected back onto the unit sphere. When the desired resolution is reached, the triangular faces are finally transformed into pentagonal and hexagonal faces: The center points of any two adjacent triangles are joined and then projected back onto the unit sphere. The previous triangle representation is simply discarded.

Although this procedure indeed yields a quite homogeneous tessellation of the unit sphere, the resulting polytope is not regular. In fact, the largest available symmetry group is the one of the initial icosahedron. Apart from this group, there are hence no further reflection transformations that map the tessellated sphere onto itself. This in turn prevents the possibility of directly matching 3*d* histograms in the way it can be done in $\mathbb{R}^2$, namely by progressively rotating the histogram's mirror image in some suitable canonical scheme. Instead, for any individual rotation a new histogram actually needs to be computed from scratch from the EGI. This procedure is very expansive, as it amounts to processing the whole EGI (which is as complex as the given input mesh) over and over again.

To avoid this computational complexity, Sun et al. [SS97] follow a simplified strategy. They consider the directions of the histogram bins as candidates for a possible symmetry plane normal. When testing a selected bin for its degree of symmetry, the histogram is just mirrored at the plane orthogonal to the bin's face normal (by applying the corresponding Householder transformation). The bin directions of the mirrored histogram fall into corresponding bins of the original histogram. Based on this correspondence, the correlation of the set of bin values is computed.

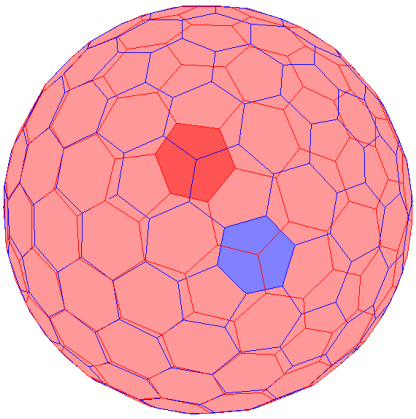

**Figure 4:** *The error that is introduced by Sun et al. [SS97]. Red edges mark the bins of the original histogram. Blue edges mark the bins after reflection on the plane corresponding to one selected face.*

This procedure is problematic for multiple reasons. First of all, reflection on the selected plane will not necessarily map the histogram frame onto itself, i. e. the two histogram frames (original and mirror) may not fit each other. This issue described in the previous paragraphs is completely disregarded. In fact, *none* of the planes considered by Sun et al. [SS97] is able to fulfill this criterion. In consequence, mirroring the histogram and transferring the histogram values to the new bins may assign normals to a cell that would not have contributed to that bin if the histogram was rather built on the mirrored EGI. Figure 4 shows that this difference can be significant for some of the bins. Red edges mark the bins of the original histogram, blue edges mark the bins after reflection. In some cases the red and blue bordered bins fit almost perfectly, while in other cases they differ a lot.

Even both extreme cases do occur. The highlighted red cell, for example, will receive no bin value from the rotated histogram, since each of the three intersecting blue bordered cells assigns its bin value to somewhere else. For the highlighted blue cell in turn the bin value will be assigned to one of the three intersecting red bordered cells. Each of these cells, however, already receives the bin value from another cell of the rotated histogram. Hence, the procedure will cause some bins to stay empty while other bins will be charged twice.

Finally, testing *all* bins for their corresponding plane's degree of symmetry in order to find the most suitable one, turns out to be still quite expensive. With increasing resolution of the tessellation, the number of bins can become very large. Moreover, computing the histogram's mirror image and mapping the mirrored bin values back to the original histogram cannot be done in linear time. It amounts to looking up every mirrored bin direction in the original histogram, which even by using a hierarchical approach will take at least $O(n \cdot \log n)$. In order to keep the computational cost low, Sun et al. [SS97] at this point consult the initially determined three principal axes and reduce the global search to a local search in the neighborhood of these directions. More precisely, for each of the three principal axes, the corresponding cell plus its five to six respectively adjacent cells are examined. Altogether, only 18–21 planes are assessed regarding their suitability for serving as a symmetry plane.

It should be noted that Sun et al. [SS97] do not clarify to which resolution their sphere is tessellated. This resolution, however, determines the angle between normals of adjacent cells and thus between the corresponding plane orientations that will be considered. Based on the figures in their article, we can only assume that Sun et al. [SS97] have used a subdivision level of 3 for their tessellation (see also the histogram shown in Figure 5, bottom). For this setting, the directions of two adjacent bins (i. e., the corresponding plane normals) differ by 7.9°, and neither smaller nor larger angles are considered, except for the other two principal axes and their local neighborhood.

The most crucial issue on Sun et al.'s approach [SS97], however, is that no kind of orientation normalization is being applied to the data, the EGI, or the orientation histogram. This directly leads to the observation that the quality of fit that can be achieved for the to-be-determined reflection plane is inevitably sensitive to the orientation of the data. In fact, the possible plane orientations are restricted to the *fixed* directions of the histogram bins. So, within the catchment area of a particular bin, the input data can be continuously rotated, while the absolute orientation of the computed reflectional plane stays fixed. That is, the plane does not rotate in alignment with the input data. Moreover, there must hence be a critical angle for which a tiny variation in the data orientation will cause the algorithm to snap to another bin (which does not even have to be a neighbor of the previous bin). The effect will be a discontinuous jump of the plane orientation by (at least) the angle described in the previous paragraph.

To treat the several issues raised by the approach due to Sun et al. [SS97], we adapted their method in three relevant points. The first of all is yet rather a minor change. In lieu of using face normals we instead refer to the surface mesh's vertex normals for building the orientation histogram. Accordingly, as corresponding individual weight we use the sum of the areas of all incident faces. The preference for vertex normals over face normals is, among others, due to the fact that for certain mesh generation methods the resulting face normals may be highly restricted in their direction. When performing Marching Cubes on a binary mask, for example, only 256 different configurations may appear, which in turn will result in only a handfull of different face normal directions. Vertex normals, on the contrary, show a much more homogeneous and smooth distribution, which should slightly improve stability.

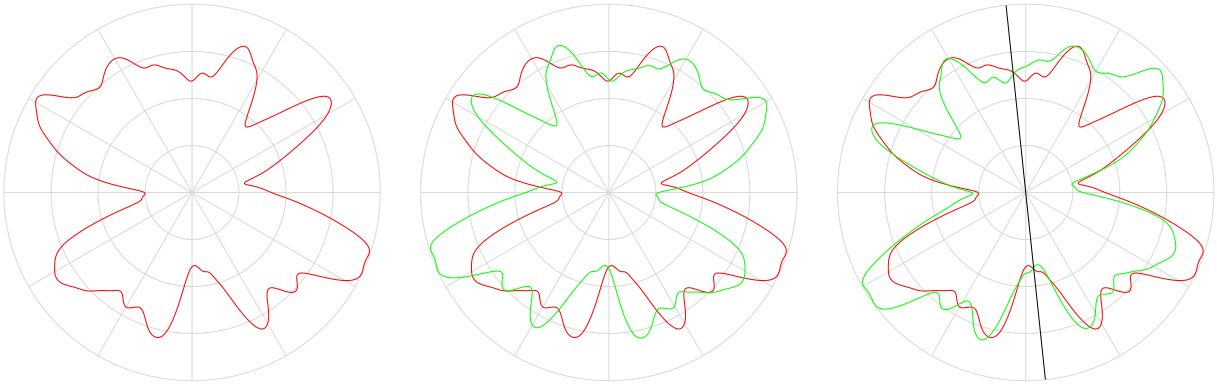

**Figure 5:** Left: 2*d orientation histogram for the bones' surface mesh shown in Figure 1*. Middle: *The same histogram and its mirror image.* Right: *Least squares matching of the histogram and its stepwise rotated mirror image. The optimum was found for a rotation of 348° = -12°, i. e. the symmetry axis is at -6°.*

As a second change, we apply an initial transformation to achieve a normalized orientation of the data. The three principal axes (which were also computed by Sun et al. [SS97], but not used in that way) constitute an orthonormal basis and induces a corresponding basis transformation that we apply to the given surface mesh. We yield a transformed representation of the data that is always in alignment with the principal axes and centered at the origin. Note that the principal component analysis (PCA) is thereby not performed on the mesh's vertices but rather on its faces (i. e. triangles). This is in particular important if the mesh does not feature a reasonably uniform resolution, as a vertex-based PCA will drift towards considerably denser parts of a mesh. In practice we use the implementation provided by the Computational Geometry Algorithms Library (CGAL) [AP09], which computes the $3 \times 3$ covariance matrix internally in closed form and not by reducing the mesh to its vertex set or by point sampling the object.

The probably most important adaption, however, is that we decided for a different way of reducing the search space. Rather than considering only a few selected orientations on a 3*d* histogram, many orientations on three 2*d* histograms are instead examined. The basic idea is that if the data is very symmetric with respect to the symmetry plane, then we should observe the same symmetry also for a projection along any support vector of that plane. (Just have a brief look at Figure 1 or 6 and imagine following a polar orbit around the mesh that stays within the plane and keeps the mesh center as focus.)

After applying the initial principal axis transformation, the data is aligned in such a way that its principal axes coincide with the *x*-, *y*-, and *z*-axis respectively. We project the mesh normals along each of this three axis (by simply omitting the corresponding coordinate for the 3*d* normal vector), yielding the orientations with respect to the *xy*-, *xz*-, and *yz*-plane. The directions of the 2*d*-vectors are sorted into bins of 1° width, thus obtaining a circular orientation histogram for each of the three principal axes. For each of these histograms the mirror image is constructed and then stepwise rotated. Rotating the mirrored histogram by 1° thereby means rotating the corresponding reflectional plane candidate by 0.5°. Note that, in contrast to 3*d* histograms, each rotation step for the 2*d* histogram takes only constant time to be computed. A least squares matching over all rotations delivers the angle with largest degree of symmetry (Figure 5). Among the three results (according to the three axes *x*, *y*, and *z*) we choose the minimal-valued one. Applying the inverse transformation of the initial orientation normalization finally yields the desired approximation of the reflectional symmetry plane.

## 4. Evaluation

For the estimation of the midsagittal plane we use a segmentation of the bony structures of head and neck. The reason is that they can be expected fairly symmetric and are not prone to deformability. Moreover, at least when dealing with CT imaging, bones are pretty easy to be segmented based on their gray value. This can be done even fully-automatically.

Once having the symmetry plane estimated, one can easily determine the sideness of the remaining anatomic structures (Figures 1 and 6). Given a corresponding surface mesh, one simply needs to perform a 3*d* orientation test for each mesh vertex, i. e. evaluate the sign of a $4 \times 4$ determinant. If all these determinants are positive, the mesh is completely to one (e. g. the left) side of the (oriented) plane. If all determinants are negative, though, the mesh is completely to the other (e. g. the right) side. Finally, if both signs show up, the mesh crosses the plane.

To evaluate the presented algorithm, a case study on 10 datasets was conducted, each case providing segmentations of head and neck structures. Besides bones, tumor and lymph nodes, also further structures were contained, such as blood

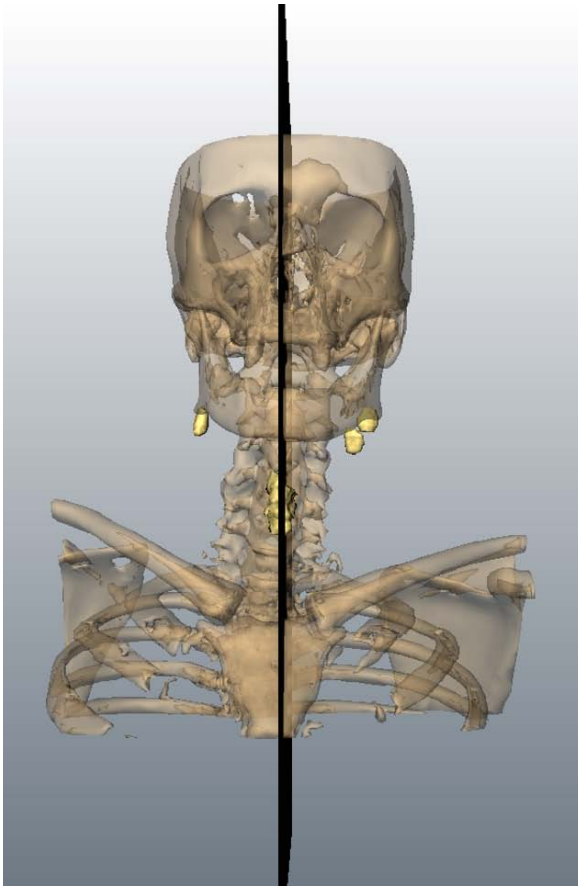

**Figure 6:** *Example for a tumor that has correctly been detected violating the midline. To enhance perception for the structures of interest, bones are visualized transparent.*

vessels, muscles, trachea or thyroid/cricoid cartilage. Due to their lack of relevance for the particular use case of TNM classification, however, these structures were not considered any further. The segmentations were created by our clinical cooperation partner, using a medical research tool for manual and semi-automatic segmentation for CT datasets [DDPS08, DTP*08]. They were provided in terms of surface triangle meshes.

For each dataset, in less than 3 seconds not only an estimation of the desired midsagittal plane could be computed, but also the sideness of tumor and lymph nodes could be determined. Altogether, 10 tumors and 134 (not necessarily affected) lymph nodes were tested for their sideness. Only three times, the sideness for a lymph node has been determined being opposite to the one stated by the medical doctor. In these cases, however, the lymph node's location with respect to the (real) midsagittal plane has actually been assessed marginal close by the medical expert.

Four times a violation of the midline by the tumor has been correctly detected (cf. Figure 6) – in the fifth case, however, the medical doctor decided in favor for the tumor still being one-sided only. (In case of a detected midline violation, though, a decision on the primary side has not been made by our application.) Two times a lymph node against medical assessment was classified as centric. For overall TNM classification, however, this did not have any effect, as centric lymph nodes are regarded ipsilateral by definition, which in turn was the correct medical conclusion in these two cases.

## 5. Discussion

We presented an approach that allows for computing an estimation of the midsagittal plane, based on a segmentation for bones of head and neck given in terms of a surface triangle mesh. The primary intention was to provide a tool for automatic sideness determination of tumor and lymph nodes in the way it is needed for tumor staging (TNM classification). Empirical tests on 10 clinical datasets have shown that the quality of the result (i. e., the midsagittal plane's computed location and orientation) in these cases was basically sufficient to either yield a correct determination of the laterality of lymph node affection or to point out a potential borderline situation in terms of a midline violation.

None of the singular cases with expert opinion differing from the computed sideness test result actually led to a modified overall assessment, subsuming the individual lateralities over all affected lymph nodes to a single compound laterality rating for TNM classification. The potential error that may have been introduced by restricting the orientation histograms to $2d$ appears negligible. Altogether, the approach has shown basically suitable for the desired purpose in medical application, in particular concerning its low required running time.

Compared to Sun et al. [SS97], our approach offers several advantages. First of all, position and orientation of the computed plane relative to the given data is insensitive to position and orientation of the latter. Moreover, the $2d$ histogram's sectional disc does not suffer from the problem of lacking reflectional self-symmetry that emerges for the sphere tessellation of a $3d$ histogram. Applying the rotation step for a $2d$ histogram takes only constant time, while not even linear time is sufficient for the $3d$ histogram of Sun et al.'s approach [SS97]. Altogether, a much greater exploration of the search space can be achieved by less computational effort. In particular, the angle between two consecutive examined planes differ by only 0.5° (note that rotating the mirrored histogram by 1° means rotating the reflectional plane candidate by 0.5°), in contrast to the $\approx 8°$ in case of Sun et al. [SS97].

However, it should be noted that the data based on which the symmetry plane was estimated, was more or less benign.

In fact, the segmented bony structure did not contain too vast artifacts that may have distracted the symmetry plane estimation significantly. In particular, the data was not truncated asymmetrically due to the bounds of image acquisition.

It remains open for future work to conduct a detailed analysis on the effective error and its quantification. In particular, a study on the average deviation of the plane computed to another plane given by a medical expert regarding their position and orientation should be conducted. Apart from that it is to examine to what extent the presented approach shows robust or sensible in the presence of rather insufficient data quality. Besides the already mentioned imaging artifacts and acquisitionally caused truncation, this also covers potential settings of atypically non-symmetric anatomy.

## References


[AP09] ALLIEZ P., PION S.: dD Principal Component Analysis. In *CGAL User and Reference Manual*, CGAL Editorial Board, (Ed.), 3.5 ed. 2009. 5

[Bru91] BRUMMER M. E.: Hough Transform Detection of the Longitudinal Fissure in Tomographic Head Images. *IEEE Trans Med Imaging 10* (1991), 74–81. 1

[DDPS08] DORNHEIM J., DORNHEIM L., PREIM B., STRAUSS G.: Modellbasierte Segmentierung von Weichgewebestrukturen in CT-Datensätzen de Halses. In *curac.08 Tagungsband* (Leipzig, September 2008), Bartz D., Bohn S., Hoffmann J., (Eds.), pp. 197–200. 6

[DTP*08] DORNHEIM J., TIETJEN C., PREIM B., HERTEL I., STRAUSS G.: Image Analysis and Visualization for the Preoperative Planning of Neck Dissections. In *5th International Conference on Computer Aided Surgery around the Head* (2008). 6

[KCD*03] KAZHDAN M., CHAZELLE B., DOBKIN D., FUNKHOUSER T., RUSINKIEWICZ S.: A Reflective Symmetry Descriptor for 3D Models. *Algorithmica 38*, 1 (2003), 201–225. 2

[KS04] KELLER Y., SHKOLNISKY Y.: An Algebraic Approach to Symmetry Detection. In *Proc IEEE ICPR* (2004), vol. 3, pp. 186–189. 2

[LCR01] LIU Y., COLLINS R., ROTHFUS W. E.: Robust Midsagittal Plane Extraction from Normal and Pathological 3-D Neuroradiology Images. *IEEE Trans Med Imaging 20*, 1 (2001), 175 – 192. 2

[MGM05] MANCAS M., GOSSELIN B., MACQ B.: Fast and Automatic Tumoral Area Localisation using Symmetry. In *Proc IEEE ICASSP* (2005), vol. 2, pp. 725– 728. 1, 2

[MIK92] MINOVIC P., ISHIKAWA S., KATO K.: Three-Dimensional Symmetry Identification Part I: Theory. *Memoirs of the Kyushu Institute of Techn. Eng. 21* (1992), 1–16. 3

[POA02] PRIMA S., OURSELIN S., AYACHE N.: Computation of the Mid-Sagittal Plane in 3-D Brain Images. *IEEE Trans Med Imaging 21*, 2 (2002), 122–138. 1, 2

[RWY95] REISFELD D., WOLFSON H., YESHURUN Y.: Context Free Attentional Operators: The Generalized Symmetry Transform. *Int J Computer Vis 14* (1995), 119–130. 2

[SS97] SUN C., SHERRAH J.: 3D Symmetry Detection using the Extended Gaussian Image. *IEEE Trans Pattern Anal Mach Intell 19*, 2 (1997), 164–168. 2, 3, 4, 5, 6

[Sun95] SUN C.: Symmetry Detection using Gradient Information. *Pattern Recognit Lett 16*, 9 (1995), 987–996. 2

[T*98] THIRION J.-P., ET AL.: Statistical Analysis of Normal and Abnormal Dissymmetry in Volumetric Medical Images. In *Proc IEEE WBIA* (1998), p. 74. 2

[TT88] TALAIRACH J., TOURNOUX P.: *Co-Planar Stereotaxic Atlas of the Human Brain: 3-Dimensional Proportional System: An Approach to Cerebral Imaging*. Thieme, New York, 1988. 1