

Extracting and Visualizing Uncertainties in Segmentations from 3D Medical Data

P. Faltin¹, K. Chaisaowong^{1,2}, T. Kraus³ and D. Merhof¹

¹Institute of Imaging and Computer Vision, RWTH Aachen University, Germany;

²King Mongkut's University of Technology North Bangkok, Thailand;

³Institute and Out-patient Clinic of Occupational Medicine, UK Aachen, Germany

Abstract

Assessing surfaces of segmentations extracted from 3D image data for medical purposes requires dedicated extraction and visualization methods. In particular, when assessing follow-up cases, the exact volume and confidence level of the segmentation surface is crucial for medical decision-making.

This paper introduces a new processing chain comprising a series of carefully selected and well-matched steps to determine and visualize a segmentation boundary. In a first step, the surface, segmentation confidence and statistical partial volume are extracted. Then, a mesh-based method is applied to determine a refined boundary of the segmented object based on these properties, whilst smoothness, confidence of the surface and partial volume are considered locally. In contrast to existing methods, the proposed approach is able to guarantee the estimated volume for the whole segmentation, which is an important prerequisite for clinical application.

Furthermore, a novel visualization method is presented which was specifically designed to simultaneously provide information about 3D morphology, confidence and possible errors. As opposed to classical visualization approaches that take advantage of color and transparency but need some geometric mapping and interpretation from the observer, the proposed scattered visualization utilizes density and scattering, which are much closer and more intuitively related to the original geometric meaning.

The presented method is particularly suitable to assess pleural thickenings from follow-up CT images, which further illustrates the potential of the proposed method.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Boundary representations, I.4.6 [Image Processing and Computer Vision]: Edge and feature detection—

1. Introduction

Pleural mesothelioma is a form of cancer, which can be caused by inhalation of asbestos fibers. The typical latency period is 20-50 years and the life expectation, after onset of the disease, is about 18 months. Therefore, patients from high-risk groups undergo a regular medical check-up, as pleural thickenings, which can act as an indicator, are observable in CT-images of the lung.

The relatively low resolution of the thickenings compared to their size makes a visual evaluation difficult, especially in the case of follow-up assessment. Due to the small and complex morphology of pleural thickenings, shown in Fig. 1(a), extracting a meaningful surface is a demanding task. The binary segmentation of thickenings can be extracted or man-

ually modified, as described in previous work [CBK*08, FNC*14]. Classical and commonly used algorithms to convert binary segmentations into surface meshes are Marching Cubes [LC87], SurfaceNets [Gib98] or more modern approaches such as convex hull based [BWC04] or quad mesh based [Nie04] approaches. They produce meshes with varying smoothness for visualization, but additional information such as target volume and segmentation confidence cannot be easily integrated. Nevertheless, this information can be included in a subsequent smoothing step. A well-known method proposed by Desbrun et al. [DMSB99] includes volume preservation as a global constraint, which does not always lead to a locally optimal smoothing, though. Bade et al. [BKP07] proposed a method, which is suitable for complex morphologies extracted from binary seg-

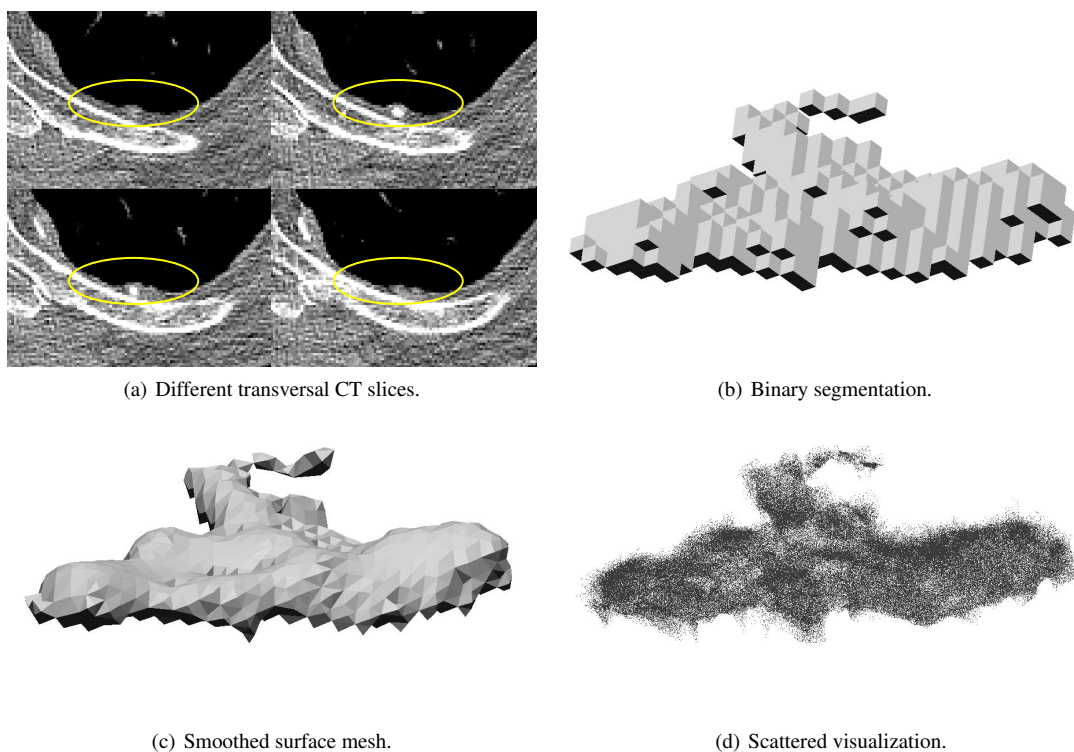


Figure 1: *Thickening with calcification located close to the spine.*

mentations. While context-aware smoothing [MGJ*11] is a method which also accounts for the alignment in the voxel raster. Essential knowledge from binary segmentations is included in both methods, but volume preservation is not guaranteed.

Laplacian-based smoothing, considering shape and volume constraints locally, to generate a smooth surface has been proposed in [FCKM14]. A method to distribute the statistically estimated partial volume is newly suggested in this work.

Classical surface mesh visualization, as shown in Fig. 1(c), provides a good impression about the presumable morphology and size of the thickening. However, the probable errors introduced in the smoothing and associated uncertainties are not visualized. This information could be included in classical mesh-based rendering using transparency and color information, but the rendering would still show a sharp surface, which gives the observer an impression of accuracy. Due to the low resolution image data and noise, this accuracy might be misleading for diagnostic purposes. The approach proposed by Kniss et al. [KVUS*05] is capable of visualizing uncertain borders between different segmentations using color and transparency but has drawbacks in visualizing uncertain surfaces not attached to other segmented objects. Therefore, we suggest a scatter plot as ren-

dering method. With a high point density it is possible to give the user an idea of the segmentation surface and still convey the uncertainty of this surface. Color can be used as an additional parameter to e.g. compare two simultaneously rendered segmentation surfaces, as shown in Fig. 8(b). Most research that addresses the visualization of point clouds covers the topic of visualizing noisy clouds as surfaces and eliminating the uncertainty in the visual context [DG04]. It is also quite common to reduce the number of points [PGK02] while preserving the visual quality. Also down- and up-sampling strategies [ABCO*01] are mainly used to obtain noise-free representations with uniform density. In contrast, we explicitly utilize noise and irregular density to represent uncertainty. An approach addressing uncertainty by point clouds was suggested by Grigoryan and Rheingans [GR04]. In contrast to our approach, this method does not address the problem of estimating a surface with correct volume and is not capable to adjust the density of the point surface.

In this work, we introduce a complete processing chain to visualize medical segmentations, including their local confidence and possible local errors. First, the partial volume (PV) of the whole thickening is statistically estimated. In the second step, this information is combined with the initial segmentation and its confidence to refine its surface. Finally, we propose a scattered visualization to display the results

and possible uncertainty of the surface. For the case of pleural thickenings, we illustrate the potential of this method by applying it to the surface part facing the pulmonary tissue.

In summary, the new contributions of this work are:

- a method to extract PV,
- a method to distribute the PV in an optimized manner,
- the possibility to adapt density and to introduce orthogonal scattering in a point-based surface visualization simultaneously
- and finally a combination of these techniques to provide a new way of assessing segmentations with uncertainty.

2. Methods

This section describes our process chain, which is divided into partial volume estimation, surface extraction and visualization.

2.1. Partial Volume Estimation

The image is defined on a discrete grid $\mathbf{R} = \{\mathbf{r} = (x, y, z)^T : 1 \leq x \leq N_x, 1 \leq y \leq N_y, 1 \leq z \leq N_z\}$. For each grid position \mathbf{r} the attenuation value in Hounsfield-units $a(\mathbf{r})$ is available. The set of segmented grid positions, belonging to a thickening, is denoted by $\mathbf{R}_S \subseteq \mathbf{R}$. Especially, when manually correcting small but complex binary segmentations, the medical expert might decide for each single voxel whether it belongs to the segmented object or not. Based on an idea proposed by Kuhnigk et al. [KDB*06] the partial volume is calculated in a region $\mathbf{R}_{S,pv} \subseteq \mathbf{R}$, defined on the image grid, which surrounds the discrete segmentation surface. This region $\mathbf{R}_{S,pv}$ is the difference of a dilated version $\mathbf{R}_{S,dilated}$ and eroded version $\mathbf{R}_{S,eroded}$ of the original thickening mask \mathbf{R}_S . We chose a sphere of radius two voxels as structuring element for the morphological operations. In contrast to the calculation of the volume, based on the average attenuation of the tissue types [KDB*06], we calculate the statistically expected volume. Our method considers the attenuation distribution for the different tissue classes. We assume that the PV included in each voxel \mathbf{r} is statistically independent from its neighbors' PV and calculate the expected volume for a single voxel using

$$E[V(S)|A = a(\mathbf{r})] = \sum_{s \in \{0,1\}} V(s) \cdot P(S = s|A = a(\mathbf{r})), \quad (1)$$

with

$$V(s) = \begin{cases} 0 & s = 0 \\ 1 & s = 1 \end{cases}. \quad (2)$$

The total expected volume of a thickening is therefore

$$V_{PV} = \sum_{\mathbf{r} \in \mathbf{R}_{S,pv}} E[V(S)|A = a(\mathbf{r})] + V(1) |\mathbf{R}_{S,eroded}|, \quad (3)$$

where $|\mathbf{R}_{S,eroded}|$ is the number of voxels in the eroded mask. The conditional probability distribution is $P(S|A) =$

$\frac{P(S \cap A)}{P(A)}$, where $P(A \cap S)$ and $P(A)$ are extracted from manual reference segmentations of pleural thickenings and lung tissue.

2.2. Surface Extraction

This section covers the complete mesh-based surface extraction. Input data are the binary segmentation, the estimated volume from Sec. 2.1 and the image data. The generated output is a smoothed surface mesh, the surface distance to the binary segmentation and a local segmentation confidence.

2.2.1. Non-Smooth Conversion from Voxel Data to Mesh Surface

The exact voxel segmentation is chosen as a starting mesh, by applying the cuberille method [HL79]. In contrast to more sophisticated methods, this one results in a volume identical to the binary segmentation. The extracted triangular mesh is described by a set of vertex positions $\mathbf{P} = \{\mathbf{p} : \mathbf{p} \in \mathbb{R}^3\}$, a set of faces $\mathbf{F} = \{\mathbf{f} = (i, j, k)^T : 1 \leq i, j, k \leq |\mathbf{P}|\}$ and a set of edges $\mathbf{G} = \{\mathbf{g} = (i, j)^T : 1 \leq i, j \leq |\mathbf{P}|\}$, connecting the vertices. A resulting surface mesh is shown in Fig. 1(b).

Additionally, we extract information from the volume data, comparable to the approach proposed by Bade et al. [BKP07]. We use the differences of the attenuation $a(\mathbf{r})$ between the voxels separated by a face. For each vertex the average of all differences from connected faces is calculated. This term is some kind of face individual contrast between the inside and outside of the segmentation. Its absolute value $q'(i)$ is clipped to the range of $[q'_{min}, q'_{max}]$ and then linearly mapped to $\lambda(i) \in [0, \lambda_{max}]$,

$$\lambda(i) = \begin{cases} 0, & |q'(i)| < t_{low} \\ q' \cdot |q'(i)|, & t_{low} < |q'(i)| < t_{high} \\ t_{slope} \cdot t_{high}, & t_{high} < |q'(i)| \end{cases} \quad (4)$$

as a measurement of segmentation confidence.

2.2.2. Smoothing

Liu et al. [LBSP02] presented a method where smoothing is carried out for each edge $\mathbf{g} = (i, j)^T$ separately. Each time only the vertex neighbors $\mathbf{N}(i)$, $\mathbf{N}(j)$ are considered which divides the mesh into small patches $\mathbf{B}(\mathbf{g} = (i, j)^T) = \mathbf{N}(i) \cap \mathbf{N}(j)$. An exemplary patch is shown in Fig. 2. The two vertex positions $\mathbf{p}(i)$, $\mathbf{p}(j)$ of an edge \mathbf{g} are first individually adapted to smooth the associated patch and in a second step modified simultaneously to preserve the initial volume. The known equations and derivations from Liu et al. [LBSP02] are kept short for brevity. Please refer to their publication for more details.

For the first step, the energy of the local patches is de-




Figure 2: Patch $\mathbf{B}(\mathbf{g})$, defined by all neighboring vertices of the edge $\mathbf{g} = (i, j)^T$.

scribed using the umbrella operator [Tau95]

$$\begin{aligned} L(i) &= \frac{1}{|\mathbf{N}(i)|} \sum_{j \in \mathbf{N}(i)} (\mathbf{p}(j) - \mathbf{p}(i)) \\ &= \omega(i, i)\mathbf{p}(i) + \sum_{j \in \mathbf{N}(i)} \omega(i, j)\mathbf{p}(j), \end{aligned} \quad (5)$$

where $\mathbf{N}(i) \subseteq \mathbf{B}(\mathbf{g})$ are the neighbors of the vertex i inside the patch $\mathbf{B}(\mathbf{g})$. All weights $\omega(i, j)$ not defined by this equation are equal to 0. The local energy function for the initial unconstrained smoothing is

$$E^1(\mathbf{g} = (i, j)^T) = \frac{1}{2} \left(L^2(i) + L^2(j) + \sum_{b \in \mathbf{B}(\mathbf{g})} L^2(b) \right), \quad (6)$$

which can be expressed in matrix notation

$$\begin{aligned} E^1(\mathbf{g} = (i, j)^T) &= \frac{1}{2} \left\| \mathbf{k}(i)\mathbf{p}^{1T}(i) + \mathbf{k}(j)\mathbf{p}^{1T}(j) - \mathbf{Q}(\mathbf{g}) \right\|_F^2 \\ &= \frac{1}{2} \left\| \mathbf{K}(\mathbf{g})\mathbf{P}^{1T}(\mathbf{g}) - \mathbf{Q}(\mathbf{g}) \right\|_F^2, \end{aligned} \quad (7)$$

where F is the Frobenius norm. By minimizing this energy $E^1(\mathbf{g})$, the new positions of the edge vertices are given by

$$\begin{aligned} \begin{bmatrix} \mathbf{p}^1(i) & \mathbf{p}^1(j) \end{bmatrix} &= \\ \mathbf{P}^1(\mathbf{g}) &= ((\mathbf{K}^T(\mathbf{g}) * \mathbf{K}(\mathbf{g}))^{-1} (\mathbf{K}^T(\mathbf{g}) * \mathbf{Q}(\mathbf{g})))^T, \end{aligned} \quad (8)$$

where $\mathbf{K}(\mathbf{g} = (i, j)^T) = [\mathbf{k}(i) \ \mathbf{k}(j)]$. The point positions of the different steps are denoted by $\mathbf{p}^0(i)$ for the initial position and $\mathbf{p}^1(i)$ for the position after smoothing. The column vectors $\mathbf{k}(i)$ and $\mathbf{k}(j)$ are filled with the entries $\mathbf{k}_b(v) = \omega(v, b) \forall v \in \{i, j\}, b \in \mathbf{B}(\mathbf{g})$ and the rows of $\mathbf{Q}(\mathbf{g})$ are defined as $\mathbf{q}_b(\mathbf{g}) = \omega(b, i)\mathbf{p}^{0T}(i) + \omega(b, j)\mathbf{p}^{0T}(j) - L(b) \forall b \in \mathbf{B}(\mathbf{g})$.

The local volume preservation is carried out in a second step. Both vertices of the edge \mathbf{g} are moved by an identical correction vector $\mathbf{w}(\mathbf{g})$ and the resulting positions are $\mathbf{P}^2(\mathbf{g}) = \mathbf{P}^1(\mathbf{g}) + [\mathbf{w}(\mathbf{g})\mathbf{w}(\mathbf{g})]$. In this step, we also include the PV and the penalty for point distance $\left\| \mathbf{p}^0(\mathbf{g}) - \mathbf{p}^2(\mathbf{g}) \right\|$, adjusted by the parameter $\lambda(i)$, as shown in [FCKM14]. Including all criteria the optimization problem, constrained by

the volumes $\Omega(\mathbf{P})$ and $\Omega(\mathbf{g})$, is

$$\begin{aligned} \min_{\mathbf{w}(\mathbf{g})} \left\{ E^2(\mathbf{g} = (i, j)^T) \right\} = \\ \min_{\mathbf{w}(\mathbf{g})} \left\{ \frac{1}{2} \left\| \tilde{\mathbf{K}}(\mathbf{g})\mathbf{w}(\mathbf{g})^T - \tilde{\mathbf{Q}}(\mathbf{g}) \right\|_F^2 + \right. \\ \left. \frac{1}{2} \lambda(i) \left\| \mathbf{p}^1(i) + \mathbf{w}(\mathbf{g}) - \mathbf{p}^0(i) \right\|^2 + \right. \\ \left. \frac{1}{2} \lambda(j) \left\| \mathbf{p}^1(j) + \mathbf{w}(\mathbf{g}) - \mathbf{p}^0(j) \right\|^2 \right\}, \end{aligned}$$

subject to $\mathbf{n}^T(\mathbf{g}) \cdot \mathbf{w}(\mathbf{g}) = \Omega(\mathbf{P}^0(\mathbf{g})) - \Omega(\mathbf{P}^1(\mathbf{g})) + \Omega(\mathbf{g})$, (9)

where $\tilde{\mathbf{K}}(\mathbf{g}) = \mathbf{k}(i) + \mathbf{k}(j)$, $\tilde{\mathbf{Q}}(\mathbf{g}) = \mathbf{Q}(\mathbf{g}) - \mathbf{K}(\mathbf{g})\mathbf{P}^{1T}(\mathbf{g})$ and $\lambda(i)$ is the segmentation confidence from Sec. 2.2.1. The constraint is defined by the volume function $\Omega(\mathbf{P})$, which calculates the volume of the open patch and the additional local patch volume $\Omega(\mathbf{g})$, calculated from the PV as explained in Sec. 2.2.3. The vector $\mathbf{n}(\mathbf{g})$ separates the induced volume change in the current patch multiplicative from $\mathbf{w}(\mathbf{g})$. The exact definitions can be found in the paper of Liu et al. [LBSP02]. While the original algorithm is fixed to a single working point ($\lambda(i) = 0$), introducing $\lambda(i)$ as a new parameter allows choosing between shape preservation and smoothness.

The resulting correction vector for our extended minimization problem is

$$\begin{aligned} \mathbf{w}(\mathbf{g}) &= \frac{1}{\alpha(\mathbf{g})} \left(\tilde{\mathbf{Q}}^T(\mathbf{g})\tilde{\mathbf{K}}(\mathbf{g}) - \lambda(\mathbf{g}) - \beta(\mathbf{g})\mathbf{n}(\mathbf{g}) \right), \\ \alpha(\mathbf{g}) &= \tilde{\mathbf{K}}^T(\mathbf{g})\tilde{\mathbf{K}}(\mathbf{g}) + \lambda(i) + \lambda(j), \\ \beta(\mathbf{g}) &= \frac{1}{\|\mathbf{n}(\mathbf{g})\|^2} \left(\mathbf{n}^T(\mathbf{g}) \left(\tilde{\mathbf{Q}}^T(\mathbf{g})\tilde{\mathbf{K}}(\mathbf{g}) - \lambda(\mathbf{g}) \right) - \right. \\ &\quad \left. \alpha(\mathbf{g}) \left(\Omega(\mathbf{P}^0(\mathbf{g})) - \Omega(\mathbf{P}^1(\mathbf{g})) + 6 \cdot \Omega(\mathbf{g}) \right) \right), \end{aligned}$$

where $\lambda(\mathbf{g}) = \lambda(i)(\mathbf{p}^1(i) - \mathbf{p}^0(i)) + \lambda(j)(\mathbf{p}^1(j) - \mathbf{p}^0(j))$.

Side note: Beside the conceptual changes in $E^2(\mathbf{g})$ and the following results, the set $\mathbf{B}(\mathbf{g})$ is defined slightly different in the publication of Liu et al. [LBSP02].

2.2.3. Initial Partial Volume Distribution

For a simple distribution of the PV V_{PV} , as derived in Sec. 2.1, it can be divided by the number of patches $|\mathbf{G}|$ to obtain additional local patch volume $\Omega(\mathbf{g}) = \frac{\Omega_{PV}}{|\mathbf{G}|}$, where $\Omega_{PV} = 6 \cdot V_{PV} - V_{\text{voxel}} \cdot |\mathbf{R}_S|$ and $|\mathbf{R}_S|$ is the number of segmented voxels. The factor of 6 is caused by the definition of the volume function $\Omega(\mathbf{P})$ [LBSP02].

What we alternatively suggest is distributing the PV, while optimizing the sum of the energy $E^2(\mathbf{g})$ for all patches

$$\begin{aligned} \min_{\Omega_g, \forall \mathbf{g} \in \mathbf{G}} \left\{ \sum_{\mathbf{g} \in \mathbf{G}} E^2(\mathbf{g}) \right\}, \\ \text{subject to } \sum_{\mathbf{g} \in \mathbf{G}} \Omega(\mathbf{g}) = \Omega_{PV}. \end{aligned} \quad (10)$$

This constrained optimization problem can be analytically solved using the Lagrange multiplier method. The resulting distributed volume per patch $\mathbf{B}(\mathbf{g})$ is given by

$$\Omega(\mathbf{g}) = \tilde{b}(\mathbf{g})(\eta - c(\mathbf{g})), \quad (11)$$

with

$$\eta(\mathbf{g}) = \left(\sum_{\mathbf{g} \in \mathbf{G}} \tilde{b}(\mathbf{g}) \right)^{-1} \left(\Omega_{PV} + \sum_{\mathbf{g} \in \mathbf{G}} \tilde{b}(\mathbf{g})c(\mathbf{g}) \right),$$

$$c(\mathbf{g}) = \sum_{n=1}^3 \sum_{\kappa} \mathbf{a}_{n,\kappa}(\mathbf{g}) \mathbf{b}_{n,\kappa}(\mathbf{g}),$$

$$\tilde{b}(\mathbf{g}) = \left(\sum_{n=1}^3 \sum_{\kappa} \mathbf{b}_{n,\kappa}^2(\mathbf{g}) \right)^{-1},$$

$$\mathbf{a}_1(\mathbf{g}) = \tilde{\mathbf{w}}(\mathbf{g}) \tilde{\mathbf{K}}^T(\mathbf{g}) - \tilde{\mathbf{Q}}^T(\mathbf{g}), \quad \mathbf{b}_1(\mathbf{g}) = 6 \frac{\mathbf{n}(\mathbf{g}) \tilde{\mathbf{K}}^T(\mathbf{g})}{\|\mathbf{n}(\mathbf{g})\|^2},$$

$$\mathbf{a}_2(\mathbf{g}) = \lambda(i) \left(\mathbf{p}^d(i) + \tilde{\mathbf{w}}(\mathbf{g}) \right), \quad \mathbf{b}_2(\mathbf{g}) = 6 \frac{\lambda(i)}{\|\mathbf{n}(\mathbf{g})\|^2} \mathbf{n}(\mathbf{g}),$$

$$\mathbf{a}_3(\mathbf{g}) = \lambda(j) \left(\mathbf{p}^d(j) + \tilde{\mathbf{w}}(\mathbf{g}) \right), \quad \mathbf{b}_3(\mathbf{g}) = 6 \frac{\lambda(j)}{\|\mathbf{n}(\mathbf{g})\|^2} \mathbf{n}(\mathbf{g}),$$

where $\mathbf{p}^d(i) = \mathbf{p}^1(i) - \mathbf{p}^0(i)$ and κ represents the element index inside a vector. The term $\tilde{\mathbf{w}}(\mathbf{g})$ is given by $\mathbf{w}(\mathbf{g})$ without the term $6 \cdot \Omega(\mathbf{g})$. The volume distribution $\Omega(\mathbf{g})$ considers volume constraint, smoothness and shape. Theoretically this procedure has to be repeated after modifying a single patch. Practically, an initial volume distribution, before modifying the mesh, is sufficient.




Figure 3: A surface of a sphere rendered as scattered visualization. The adjustable parameters are demonstrated by increasing confidence λ from left to right and increasing point distance d from bottom to top.

2.3. Scattered Visualization

The scattered visualization method takes the smoothed surface with the vertices $\mathbf{P}^2(\mathbf{g})$ as an input. The scattered sur-

face is generated in three steps. First, the input mesh is subdivided into a high number of triangles to give sufficient candidates to render points of the surface. Second, a point selection, based on a dart throwing method [CJW*09], is used to select a given number m of points in a stochastic process. This step also includes the confidence $\lambda(i)$ to adjust the density of the point-based surface. In the third step, the point distance $d(i) = \|\mathbf{p}^0(i) - \mathbf{p}^2(i)\|$ between original and smoothed segmentation is included in the visualization as scatter orthogonal to the surface. A resulting scattered visualization of a sphere with varying parameters is shown in Fig. 3.

2.3.1. Mesh Subdivision

The subdivision process is not meant to refine the detail level of the segmentation surface, as the extractable details are strongly limited by the input image resolution. Instead, it is used to generate a dense set of candidates for the scattered visualization. Therefore, a smooth refinement which modifies the subdivided mesh is not desired. Therefore, bisecting the triangle edges and creating four new triangles from each input triangle, as suggested by Cline et al. [CJW*09], is sufficient. The process is repeated, until a predefined maximum permitted triangle area a_{\max} is reached. The new triangles $\mathbf{f}' = (i', j', k')^T \in \mathbf{F}'$ have the centroids $\mathbf{c}'(\mathbf{f}')$ and the area $a'(\mathbf{f}')$. The point distances $d(i)$, arising from the smoothing in Sec. 2.2 and the confidences $\lambda(i)$ from Sec. 2.2.1 are mapped to the new faces \mathbf{f}' and given by $d'(\mathbf{f}')$ and $\lambda'(\mathbf{f}')$. For smooth transitions at the triangle edges, we use an inverse distance weighting [She68] for interpolation. E.g. the confidence for each new face \mathbf{f}' is given by

$$\lambda'(\mathbf{f}') = \frac{w(\mathbf{f}', i)\lambda(i) + w(\mathbf{f}', j)\lambda(j) + w(\mathbf{f}', k)\lambda(k)}{w(\mathbf{f}', i) + w(\mathbf{f}', j) + w(\mathbf{f}', k)}, \quad (12)$$

with $w(\mathbf{f}', i) = \frac{1}{(\mathbf{p}(i) - \mathbf{c}(\mathbf{f}'))^2}$ and i, j, k are the vertices of the original triangle before subdivision.

2.3.2. Point-based Surface

A scattered visualization requires an adjustable density of points, which shall not depend on the underlying triangle size and shape. We utilize a dart throwing algorithm [CJW*09] and extend it by including the certainty $\lambda'(\mathbf{f}')$. Instead of using the area $a'(\mathbf{f}')$, we use the product $b'(\mathbf{f}') = a'(\mathbf{f}') \cdot \lambda'(\mathbf{f}')$ as the selection probability. Thereby, the number of points at the boundary region is scaled inversely with its segmentation confidence. Then we follow the original algorithm [CJW*09]: Each candidate \mathbf{f}' is sort into a bin $B(\mathbf{f}')$, where $B(\mathbf{f}') = \left\lfloor \log_2 \left(\frac{\max(b'(\mathbf{f}'))}{b'(\mathbf{f}')} \right) \right\rfloor$. A triangle is chosen by selecting a bin with probability, proportional to the total sum of $b'(\mathbf{f}')$. Then, a triangle is chosen within the bin and accepted with probability $\frac{b'(\mathbf{f}')}{\max(b'(\mathbf{f}'))}$. Otherwise the triangle is rejected and a new one is chosen from the same bin. This is repeated until a triangle is accepted. Due




Figure 4: Roughness of sphere surface shows effect of different volume distribution methods.

to the logarithmic binning the acceptance rate is at least fifty percent and the triangle selection will run in constant time on average [CJW*09]. The point selection is repeated until the desired number m of points has been reached.

The certainty $\lambda'(\mathbf{f}')$ just describes the relative density and does not influence the total density. The latter can only be adjusted by the number of points m . It is necessary to produce identical densities for different visualized objects of the same scenario; a scenario is given by the relevant tissue contrast in the image and the resolution of the object. Therefore, we choose the number of points m as the sum of all selection probabilities

$$m = \left\lceil v \sum_{\mathbf{f}' \in \mathbf{F}'} b'(\mathbf{f}') \right\rceil, \quad (13)$$

where v is a parameter, which should be chosen depending on the scenario. In our scenario of pleural thickenings, we chose $v = 2000$.

2.3.3. Orthogonal Scattering

In the final step, each resulting point from Sec. 2.3.2 is equally distributed orthogonal to the surface. The interval is $[-d'(\mathbf{f}'), d'(\mathbf{f}')]$, with the interpolated point distance $d'(\mathbf{f}')$ from Sec. 2.2.1. The orthogonal direction is given by the surface normal. For a visually smooth transition the inverse distance weighting [She68] is applied to interpolate the normals and the point distances $d'(\mathbf{f}')$, analogous to the confidence interpolation in step 2.3.1.

3. Results

The result section is split into three parts. First, the surface extraction is evaluated regarding the adjustability between smoothness and shape preservation. The second part analyzes the visualization step for a single object. Beside qualitative discussions of the strengths and weaknesses, a quantitative analysis of the uniformity of point density is given.

Finally, a short and qualitative analysis illustrates the scenario of comparing follow-up segmentations.

3.1. Surface Extraction

The influence of the segmentation confidence was assessed using a sphere with radius of 2 cm consisting of 4098 vertices. Its vertex positions $\mathbf{p}(i)$ are disturbed by white additive noise in the interval of $[-0.1, 0.1]$ cm, resulting in a mesh as shown in Fig. 4(a). The smoothing was carried out with parameter $\lambda(i) \in [0, 16]$ (constant for all vertices) and for 1 and 5 iterations. Additionally, the simple method of distributing the volume equally to all patches is compared to the optimized volume distribution described by Eq. 10. In Fig. 5, the resulting average vertex distance to the noisy (5(a)) and original (5(b)) sphere is plotted versus the mesh roughness, given by the average length of the Geometric Laplacian [Lav07]. The dashed lines are benchmark values, representing the average distance between the vertices of the noisy and the initial sphere (horizontal line) and the average length of the Geometric Laplacian [Lav07] for the initial (left line) and the noisy (right line) sphere. Since there is no confidence parameter in the original algorithm [LBSP02] its only possible working points are given with $\lambda(i) = 0$ and marked by circles in Fig. 5. Using the optimized volume distribution one iteration is sufficient for results close to the convergence area, which is otherwise only reached with approximately five iterations.

The effect of the proposed volume distribution method in Sec. 2.2.3 can also be visually observed. The noisy sphere of Fig. 4(a) is smoothed with $\lambda(i) = 0$ for all patches. The simple method, with an identical distribution for every patch, results in an irregular shape, shown in Fig. 4(b). Whereas the optimized method recovers the sphere with a more regular shape, recognizable in Fig. 4(c). An exemplary smoothing result for a thickening surface with 5 iterations, using the optimized volume distribution method, $q'_{\min} = -1024$,

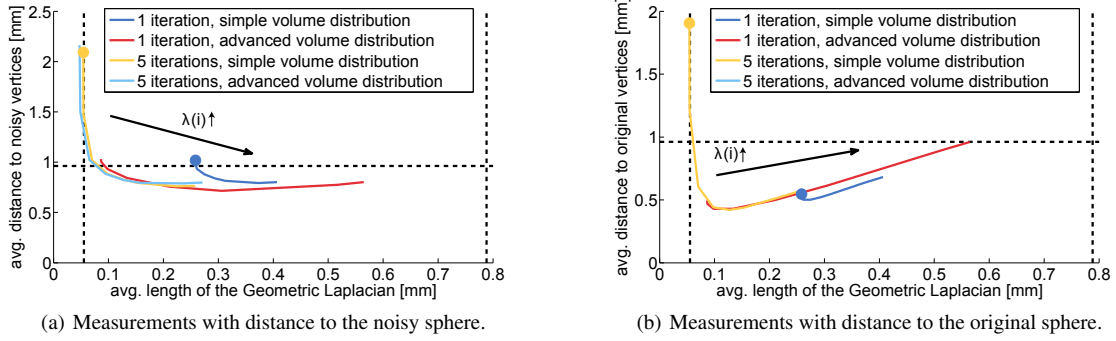


Figure 5: Roughness is plotted versus shape error. The working point is adjustable by the parameter $\lambda(i)$.

$q'_{\max} = 200$ and $\lambda_{\max} = 1$ is shown in Fig. 1(c). In all experiments the exact given volume is preserved.

3.2. Scattered Visualization

An exemplary pleural thickening was rendered in Fig. 1(d). The different light and dark spots are not caused by coloring or shading, but by varying density of the points, which represents the confidence $\lambda'(\mathbf{f}')$. Uncertainty and shape are visualized in a comprehensible way. However, in static 2D views, the shape of the object cannot be easily observed without additional help (e.g. a classical mesh-based rendering). Therefore, it can be interpreted as a more error-oriented, instead of shape-oriented, visualization method.

The point density is an important indicator for the confidence. Therefore an evaluation of the uniformity was carried out. A random mesh with size of 10×10 pixels, as shown in Fig. 6, was created and visualized with error $d = 0$, constant confidence $\lambda'(\mathbf{f}') = 1, \forall \mathbf{f}' \in \mathbf{F}'$, varying maximal area a_{\max} and varying number of points m . Then, the resulting numbers of points per pixel were analyzed.

The average number of points per pixel \bar{v}_{pixel} with error bars indicating the standard deviation is shown in Fig. 7(a). The variation coefficient of the point density $\frac{\sigma(\bar{v}_{\text{pixel}})}{\bar{v}_{\text{pixel}}}$ represents the relative variation and is shown in Fig. 7(b). As the algorithm always terminates when the desired number of point v has been reached, the average number in all visualized pixels is always exactly achieved. The density deviation between the simulated 10×10 pixels is given by the error bars, which shows that the absolute variance is slightly increasing for larger numbers of points. Fig. 7(b) shows that the observable density variation is decreasing for an increasing number of points and approaching a minimal value of about 0.05. The minimal area a_{\max} has no observable influence on the point density. However, the minimal area is limited, because it must be small enough to generate a sufficient number of point candidates. Otherwise, the algorithm cannot terminate in step 2.3.2. Due to the stronger variation





Figure 6: Plane with irregular shaped triangles to evaluate the point density.

for a small number of points, a visual inference of the actual confidence from the point density is not exactly possible in these cases. However, a minimal number of points are required anyway, to visually convey the shape of the object.

3.3. Visual Comparison and Growth Estimation


An important aspect of follow-up assessment of pleural thickenings is the growth estimation. An overlay of mesh based representation, as shown in Fig. 8(a), might lead to growth estimation in regions with very uncertain segmentation boundaries. The scattered visualization of this overlay, in Fig. 8(b), allows a more informed decision about growth in different regions.




(a) Average number and standard deviation of points in a pixel, depending on total number of points.

(b) Standard deviation normalized to the number of average points, depending on total number of points for different maximum area of the initial triangulation.

Figure 7: Reproducibility of point density.



(a) Overlay of mesh based visualization.



(b) Overlay of scattered visualization.

Figure 8: Follow-up assessment of pleural thickenings from two different points in time. The first point in time is shown in orange and the second point in time in blue.

3.4. Applications

The method was mainly designed for the growth assessment of pleural thickenings. However, it is applicable in a variety of contexts, where uncertain boundaries have to be visualized. One of the main advantages is the exact volume representation. Therefore, it is especially useful to visually assess thin objects or objects with complex morphologies.

One useful application might be the assessment of gray matter, which is an important indicator in the field of neuroscience. The available image data has typically a relatively low resolution, compared to the complex morphology of the gyri and sulci. Another interesting field of application could be the visualization of cartilage, as images typically suffer from low image contrast. Especially in the case of articular cartilage of the knee, the thickness is an important criterion to detect or quantify damage.

4. Conclusions and Outlook

We presented a comprehensive tool to extract, process and visualize shape and uncertainty of 3D objects. The presented smoothing method is vertex-wise adjustable between shape preservation and smoothness, which is used to discriminate voxels segmented with varying confidence. Furthermore, the optimized distribution of the volume among the surface results in a fast converging smoothing process and in a regular mesh shape. The visualization part can successfully convey the knowledge about shape, confidence and modifications caused by the smoothing. Limitations of the uniformity of point density only occur for a very small number of points, which is not useful to visualize a closed surface anyways. Especially when comparing follow-up images, the additional information is of great benefit to judge changes between the

two images. Nevertheless, it should be noted that this representation is less suitable for a 2D static visualization.

For the future a better interpretation of the 3D shape in the scatter plot is desirable, by e.g. superimposing additional geometric information. Another aspect is the exact visualized volume, which is hard to evaluate quantitatively, but might be influenced by the inverse distance weighting.

References

- [ABCO*01] ALEXA M., BEHR J., COHEN-OR D., FLEISHMAN S., LEVIN D., SILVA C.: Point set surfaces. In *Visualization, 2001. VIS '01. Proceedings* (Oct 2001), pp. 21–29, 537. 2
- [BKP07] BADE R., KONRAD O., PREIM B.: Reducing artifacts in surface meshes extracted from binary volumes. *Journal of WSCG 15*, 1-3 (2007), 67–74. 1, 3
- [BWC04] BHANIRAMKA P., WENGER R., CRAWFIS R.: Isosurface construction in any dimension using convex hulls. *IEEE TVCG 10*, 2 (2004), 130–141. 1
- [CBK*08] CHAISAOWONG K., BROSS B., KNEPPER A., KRAUS T., AACH T.: Detection and follow-up assessment of pleural thickenings from 3D CT data. In *Proc. ECTI* (2008), vol. 1, pp. 489 – 492. 1
- [CJW*09] CLINE D., JESCHKE S., WHITE K., RAZDAN A., WONKA P.: Dart throwing on surfaces. In *Computer Graphics Forum* (2009), vol. 28(4), Wiley Online Library, pp. 1217–1226. 5, 6
- [DG04] DEY T. K., GOSWAMI S.: Provable surface reconstruction from noisy samples. In *Proceedings of the twentieth annual symposium on Computational geometry* (2004), ACM, pp. 330–339. 2
- [DMSB99] DESBRUN M., MEYER M., SCHRÖDER P., BARR A. H.: Implicit fairing of irregular meshes using diffusion and curvature flow. In *Proc. SIGGRAPH* (1999), ACM, pp. 317–324. 1
- [FCKM14] FALTIN P., CHAISAOWONG K., KRAUS T., MERHOF D.: Consistent follow-up segmentation of pleural thickenings. In *Proc. IEEE International Symposium on Biomedical Imaging (IEEE ISBI)* (2014), pp. 1368–1372. 2, 4
- [FNC*14] FALTIN P., NGUYEN P.-A., CHAISAOWONG K., KRAUS T., MERHOF D.: Interactive 3D segmentation of pleural thickenings simultaneously at different points of time using graph cut. In *Bildverarbeitung für die Medizin 2014* (2014), Springer, pp. 378–383. 1
- [Gib98] GIBSON S.: Constrained elastic surface nets: Generating smooth surfaces from binary segmented data. *MICCAI* (1998), 888–898. 1
- [GR04] GRIGORYAN G., RHEINGANS P.: Point-based probabilistic surfaces to show surface uncertainty. *Visualization and Computer Graphics, IEEE Transactions on 10*, 5 (Sept 2004), 564–573. doi:10.1109/TVCG.2004.30. 2
- [HL79] HERMAN G. T., LIU H. K.: Three-dimensional display of human organs from computed tomograms. *Computer graphics and image processing 9*, 1 (1979), 1–21. 3
- [KDB*06] KUHNIGK J., DICKEN V., BORNEMANN L., BAKAI A., WORMANN S., KRASS S., PEITGEN H.: Morphological segmentation and partial volume analysis for volumetry of solid pulmonary lesions in thoracic CT scans. *IEEE TMI 25*, 4 (2006), 417–434. 3
- [KVUS*05] KNISS J., VAN UITERT R., STEPHENS A., LI G.-S., TASDIZEN T., HANSEN C.: Statistically quantitative volume visualization. In *Visualization, 2005. VIS 05. IEEE* (Oct 2005), pp. 287–294. doi:10.1109/VISUAL.2005.1532807. 2
- [Lav07] LAVOUÉ G.: A roughness measure for 3D mesh visual masking. In *Proc. APGV* (2007), ACM, pp. 57–60. 6
- [LBSP02] LIU X., BAO H., SHUM H.-Y., PENG Q.: A novel volume constrained smoothing method for meshes. *Graphical Models 64*, 3 (2002), 169–182. 3, 4, 6
- [LC87] LORENSEN W. E., CLINE H. E.: Marching cubes: A high resolution 3D surface construction algorithm. In *ACM Siggraph Computer Graphics* (1987), vol. 21(4), ACM, pp. 163–169. 1
- [MGJ*11] MOENCH T., GASTEIGER R., JANIGA G., THEISEL H., PREIM B.: Context-aware mesh smoothing for biomedical applications. *Computers & Graphics 35*, 4 (2011), 755–767. 2
- [Nie04] NIELSON G. M.: Dual marching cubes. In *Proc. IEEE VIS* (2004), pp. 489–496. 1
- [PGK02] PAULY M., GROSS M., KOBELT L. P.: Efficient simplification of point-sampled surfaces. In *Proceedings of the conference on Visualization'02* (2002), IEEE Computer Society, pp. 163–170. 2
- [She68] SHEPARD D.: A two-dimensional interpolation function for irregularly-spaced data. In *Proceedings of the 1968 23rd ACM national conference* (1968), ACM, pp. 517–524. 5, 6
- [Tau95] TAUBIN G.: A signal processing approach to fair surface design. In *Proc. SIGGRAPH* (1995), ACM, pp. 351–358. 4