Isosurface Reconstruction with Topology Control

Stephan Bischoff   Leif P. Kobbelt

Lehrstuhl für Informatik VIII, RWTH Aachen[*]

Abstract

Extracting isosurfaces from volumetric datasets is an essential step for indirect volume rendering algorithms. For physically measured data like it is used, e.g. in medical imaging applications one often introduces topological errors such as small handles that stem from measurement inaccuracy and cavities that are generated by tight folds of an organ. During isosurface extraction these measurement errors result in a surface whose genus is much higher than that of the actual surface. In many cases however, the topological type of the object under consideration is known beforehand, e.g., the cortex of a human brain is always homeomorphic to a sphere. By using topology preserving morphological operators we can exploit this knowledge to gradually dilate an initial set of voxels with correct topology until it fits the target isosurface. This approach avoids the formation of handles and cavities and guarantees a topologically correct reconstruction of the object's surface.

Keywords: isosurface extraction, digital topology, discrete deformation retraction, topology preservation

1   Introduction

Volumetric datasets are usually given by scalar values (= gray values, density values) being assigned to the vertices of a regular voxel grid. In order to effectively visualize and process the geometric information which is implicitly represented by a volumetric dataset, we often have to extract explicit surface information.

While the spatial topology of the regular voxel grid is very simple, the manifold topology of the extracted isosurfaces can be very complex. In practice it turns out that most of the topological complexity of isosurfaces (e.g., small handles and channels) is due to measurement inaccuracies and noise in the scalar data. For example, when extracting the cortex of a human brain (which is known to be homeomorphic to a sphere) we often find erroneous channels and disconnected components. These effects are caused by the fact that the measuring resolution of a CT or MRT scanner

it usually not fine enough to properly distinguish the tissue of nearby sulci and gyri (ridges).

The goal of this paper is to derive an algorithm that is able to extract isosurfaces with controlled topologies. The idea is to find an initial estimate of the final surface which can be guaranteed to have a desired topology. Then a topology preserving growing scheme is applied that morphs the initial estimate into the final result. The topology preservation will prevent the generation of handles such that the final isosurface might touch itself but without actually changing its manifold topology.

There are two fundamentally different concepts for accessing and modifying isosurface information represented by a volumetric dataset.

The first one is based on the pure topology of the underlying voxel grid and mostly ignores the scalar data associated with the voxels. Since the neighborhood of each voxel is regular we do not have to consider topological special cases. The well-known morphological operators are an example for operators which use spatial grid neighborhoods to determine the status of a voxel. In our algorithm we use morphological operations for the incremental growing procedure to guarantee that the extracted surface does not change its topology – even if the surface touches itself during the growing process.

Notice that the neighborhood relation in the voxel grid is defined in terms of Euclidian distances while the neighborhood relation within an isosurface is defined in terms of geodesic distance. This makes it much more difficult to detect topology changes on manifolds since geometrically nearby surface points can have a large geodesic distance.

The second concept considers the volumetric dataset as a scalar-valued function defined in continuous 3-space by interpolating the grid values. Through the implicit function theorem, this scalar field defines the local geometry of all its isosurfaces. Computing the derivative of the scalar field gives us the gradient vector which is the normal vector to the isosurface. The gradient vector is used in our algorithm to determine the direction in which the growing procedure should continue. The gradient sometimes appears in disguised form, i.e., when we morph an isosurface

$$S(v) \;=\; \{(x,y,z)|f(x,y,z)=v\}$$

into an isosurface $S(v \pm \varepsilon)$ we, in principle, shift all surface points in the direction of the gradient.

In our surface extraction algorithm we combine the two concepts by exploiting the advantages of both. We grow the initial surface by adding one voxel at a time. The scalar field data and its gradient determine *where* to add the next voxel while the topological properties of a morphological operator determine *if* a specific voxel can be added without changing the resulting surface topology.

1.1 Overview

Our approach of inflating an isosurface is governed by two principles: From a topologist's point of view, we have to guarantee that the genus of the surface does not change during inflation — a behavior that can be described conveniently by the concept of deformation retractions known from algebraic topology [6]. The goal of Section 2 is to carry over this notion from the continuous case to the discrete setting of voxel spaces. A physician on the other hand might rather be interested in the underlying gradient forces that control the expansion rate of the surface. Modeling this in the discrete setting will be the topic of Section 3. In Section 4 we demonstrate our topology preserving inflation process by applying it to the problem of reconstructing the surface of a brain from a set of MRT scans.

1.2 Related Work

Topology: Discrete topological spaces where studied since the beginning of the last century [15]. It was, however, the advent of digital computers and their inherent inability to handle continuous data that intensified the interest in *digital topology* [17]. Especially in the context of image processing numerous approaches have been presented to define a digital analog of the topology of the Euclidean space. An overview can be found e.g. in [8]. In particular the problem of characterizing *simple points*, i.e. points that can be removed without changing an objects' topology, has received a great deal of attention, see e.g. [18, 12, 4, 11].
**Medical imaging:** The reconstruction of surfaces from medical image data has received a great deal of attention in recent years. Morphological operators were applied to segment the voxel data [7, 21, 20]. In order to get a topologically correct reconstruction of the surface of an organ, most methods start out with a surface of known topology and then shrink/inflate it to match a given shape [3, 14, 2]. Kriegeskorte and Goebel [9] describe an algorithm that is very similar to ours. However, they use a different topological framework and also they do not account for gradient information in the data. As a consequence they have less control on the optimal placement of the cuts that remove erroneous handles.

Surface reconstruction: Meshes that are reconstructed from range data by volumetric methods often suffer from topological noise due to mis-registrations of the scans [1]. Guskov et al. [5] propose a local wave front traversal algorithm that identifies and removes the handles from the mesh. This approach differs from ours in that it removes the handles after extraction of the surface while our approach works directly on the voxel set.

2 Topology

2.1 Topological background

Our basic elements $\mathbf{s} \in \mathbb{Z}^3$ are called *voxels*. Voxels might either be thought of as uniform sample points of a scalar field or as unit cubic cells centered on the integer grid. Subsets $S \subset \mathbb{Z}^3$ are called *voxel sets*. We will assume that all voxel sets are bounded, i.e. $|S| < \infty$. Voxels $\mathbf{s} \in S$ are called *solid*, voxels in $\mathbb{Z}^3 \setminus S$ are called *empty*. Two solid voxels $\mathbf{s}, \mathbf{t} \in S$ are called *neighbors*, $\mathbf{s} \sim_s \mathbf{t}$, if $||\mathbf{s} - \mathbf{t}||_\infty = 1$. Two empty voxels $\mathbf{a}, \mathbf{b} \in \mathbb{Z}^3 \setminus S$ are called neighbors, $\mathbf{a} \sim_e \mathbf{b}$, if $||\mathbf{a} - \mathbf{b}||_1 = 1$. As each solid voxel can have up to 26 neighbors, $\sim_s$ is often called the 26-neighborhood relation. Analogously $\sim_e$ is called the 6-neighborhood relation (Figure 1). It is equally possible to consider the dual case, namely to use the 6-neighborhood relation for solid voxels and the 26-neighborhood relation for the empty voxels. However, for the sake of the simplicity of the exposition, we stick to the first definition.
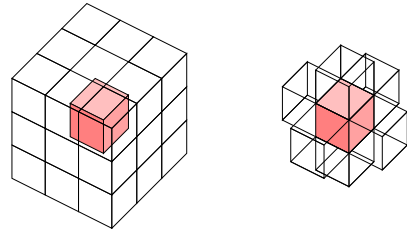
Figure 1. Different voxel neighborhood relations. Left: Vertex- or 26-neighborhood. Right: Face- or 6-neighborhood.

Let $\varepsilon \geq 0$. We define the ε-*embedding* $\mathcal{R}_\varepsilon(\mathbf{s}) \subset \mathbb{R}^3$ of a solid voxel $\mathbf{s}$ as

$$\mathcal{R}_\varepsilon(\mathbf{s}) = \left\{ \mathbf{x} \in \mathbb{R}^3 : ||\mathbf{x} - \mathbf{s}||_\infty \leq \frac{1}{2} + \varepsilon \right\}$$

and the ε-*embedding* $\mathcal{R}_\varepsilon(S) \subset \mathbb{R}^3$ of S as the union

$$\mathcal{R}_\varepsilon(S) = \bigcup_{\mathbf{s} \in S} \mathcal{R}(\mathbf{s})$$

Intuitively this means that we center a cube of edge length $1 + 2\varepsilon$ at each solid voxel $\mathbf{s} \in S$. Note that this definition

of $\mathcal{R}_\varepsilon$ nicely matches the definition of the neighborhood relation $\sim_s$, since

$$\mathbf{s} \sim_s \mathbf{t} \Leftrightarrow \mathcal{R}_\varepsilon(\mathbf{s}) \cap \mathcal{R}_\varepsilon(\mathbf{t}) \neq \emptyset.$$

In particular this demonstrates that $\sim_s$ and $\sim_e$ are *compatible*, i.e. each vertex (edge,face) is non-ambiguously assigned to be either solid or empty.
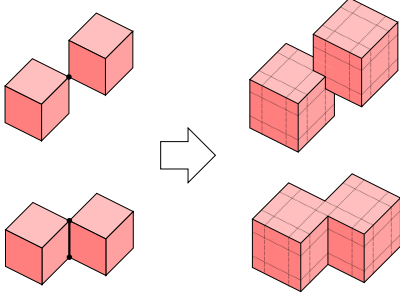
Figure 2. Singular edges and vertices are removed by scaling each voxel by a factor of $1 + \varepsilon$ about its center.

Finally, we associate a 2-manifold $\partial_\varepsilon S$ to each voxel set S by setting

$$\partial_\varepsilon S := \partial \mathcal{R}_\varepsilon(S)$$

If $0 < \varepsilon < \frac{1}{2}$ the surface $\partial_\varepsilon S$ is a regular 2-manifold, in particular, there are no singular points or singular edges (Figure 2).

2.2 Euler characteristic

In the following we associate an Euler characteristic with a voxel set S. First, we note, that a single voxel $\mathcal{R}(\mathbf{s})$ can be viewed as a finite cell complex $\mathcal{C}(\mathbf{s})$ consisting of 1 cube (3-cell), 6 quadrangles (2-cells), 12 edges (1-cells) and 8 vertices (0-cells) as depicted in Figure 3.
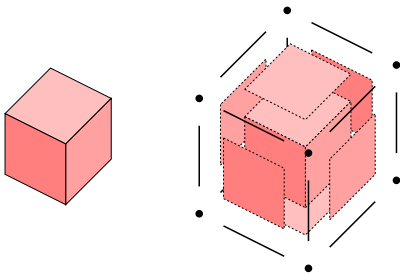


Figure 3. Decomposition of a voxel into a cell complex.

Analogously, we can can represent $\mathcal{R}(S)$ as a finite cell complex $\mathcal{C}(S)$ defined by

$$\mathcal{C}(S) = \bigcup_{\mathbf{s} \in S} \mathcal{C}(\mathbf{s})$$

It is well known that the Euler characteristic χ of a cell complex $\mathcal{C}$ is the number of even dimensional cells minus the number of odd dimensional cells [6]. Hence we define

$$\chi(S) := n_0 - n_1 + n_2 - n_3$$

where n_i is the number of i-cells ($i = 0, 1, 2, 3$) of $\mathcal{C}(S)$. Examples of this computation are shown in Figure 4.
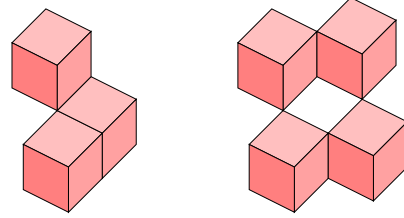
Figure 4. Computing the Euler characteristic of a voxel set. The left configuration is homeomorphic to a sphere, $\chi = 18 - 31 + 17 - 3 = 1$. The right configuration is homeomorphic to a torus, $\chi = 26 - 46 + 24 - 4 = 0$.

Implementation note

To efficiently compute the Euler characteristic of a voxel set S, we exploit the fact, that χ is additive, i.e.

$$\chi(S) = \sum_{\mathbf{a} \in \mathbb{Z}^3} \chi\left(S \cap (\mathbf{a} + [0, 1)^3)\right)$$

Each term $\chi(S \cap (\mathbf{a} + [1, 0)^3))$ in the sum above is uniquely determined by the status of the 8 voxels at positions $\mathbf{a}, \mathbf{a} + (1, 0, 0), \ldots, \mathbf{a} + (1, 1, 1)$ and can conveniently be precomputed in a lookup table of size 256 (Figure 5). This is a very important observation since it allows us to efficiently detect topology changes caused by incrementally setting one voxel to solid. Alternatively, the local configuration around a voxel can be coded in a binary decision diagram , see e.g. [16].

2.3 Minimally different voxel sets

In this section we define what it means for two voxel sets S and T to be only minimally different from each other. We will use this notion of "closeness" in the next section to carry over continuous transformations to the discrete setting of voxel spaces.

Obviously, a necessary condition for S and T to be similar to each other is that they differ by at most one voxel $\mathbf{s}$, i.e.

$$S = T \cup \{\mathbf{s}\} \quad \text{or} \quad T = S \cup \{\mathbf{s}\}.$$

Without loss of generality, we assume in the following that $S \cup \{\mathbf{s}\} = T$. By this definition alone we cannot distinguish voxel sets of different Euler characteristics. Hence
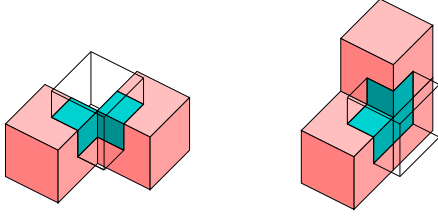
Figure 5. Computing the Euler characteristic of a voxel set can be simplified by using a lookup table. Each entry is determined by the status of a $2 \times 2 \times 2$ block of voxels. To compute the entry, we count the fractional contributions of each of the intersecting cell complexes. On the left the two solid voxels jointly contribute 1 vertex, 5 half–edges, 6 quarter–faces and 2 eigth–cubes, hence $\chi = 1 - \frac{5}{2} + \frac{6}{4} - \frac{2}{8} = -\frac{1}{4}$. For the right configuration we compute analogously $\chi = 1 - \frac{5}{2} + \frac{7}{4} - \frac{3}{8} = -\frac{1}{8}$.

we additionally require, that

$$\chi(S) = \chi(T)$$

Still this is not sufficient to prevent unwanted changes as is demonstrated in Figure 6. In this case the insertion of the center voxel at the same time removes *and* creates a handle, the change $\Delta\chi$ of the Euler characteristic is

$$\Delta\chi = 0 + 5 - 4 - 1 = 0.$$

To detect such undesired topology changes, we consider the problem on a sub-voxel scale. For this we conceptually subdivide each voxel $\mathbf{s}$ into 8 sub-voxels $\mathbf{s}_1, \ldots, \mathbf{s}_n$. Then we define S and T to be *mimimally different*, $S \approx T$ if there exists a permutation π, such that

$$
\begin{aligned}
\chi(S) &= \chi(S \cup \{\mathbf{s}_{\pi(1)}\}) \\
&= \chi(S \cup \{\mathbf{s}_{\pi(1)}, \mathbf{s}_{\pi(1)}\}) \\
&= \cdots \\
&= \chi(S \cup \{\mathbf{s}_{\pi(1)}, \cdots, \mathbf{s}_{\pi(8)}\}) \\
&= \chi(T)
\end{aligned}
$$

Implementation note Note that the choice of π is not unique, hence any such π will do. In our implementation we greedily add subvoxels $\mathbf{s}_i$ as long as χ does not change. If all subvoxels $\mathbf{s}_i$ can be added this way, we have found a suitable permutation π and hence proved that $S \approx T$.

2.4 Discrete deformation retraction

In this section we carry over the notion of a deformation retraction from the continuous case considered in algebraic topology to the discrete setting of voxel spaces.
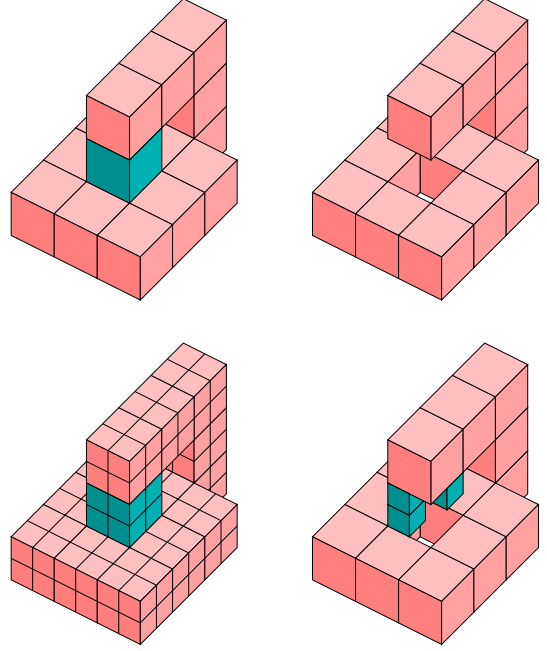
Figure 6. Discrete deformation retraction: The objects at top left and right are very similar to each other: They differ by only one voxel and they are homeomorphic. However, there is no deformation retraction that continuously transforms the left object into the right one. To model this in the discrete setting we split each voxel into 8 subvoxels (bottom left). Then we try to find a permutation π such that the successive removal of $\mathbf{s}_{\pi(1)}, \ldots, \mathbf{s}_{\pi(n)}$ does not change the Euler characteristic of the set.

A deformation retraction of a topological space X onto a subspace $Y \subset X$ is a map

$$
\begin{aligned}
f : X \times [0,1] &\rightarrow Y \\
(x, t) &\mapsto f(x, t) =: f_t(x)
\end{aligned}
$$

such that the following properties hold:

1. $f_0 \equiv id_X$, $f_1(X) = Y$ and $f_t|_Y \equiv id_Y$

2. f is continuous in x as well as in t

Intuitively this means that the space X can continuously be shrunken onto the space Y. Note that the existence of a deformation retraction implies homeomorphic equivalence of X and Y but not vice versa.

With the preparatory work of the previous sections at hand, we can now easily carry over this notion to the discrete setting: The topological spaces X and Y are replaced by voxel sets S and T and the continuous interval $[0, 1]$ is replaced by a discrete sequence of points in time numbered from 0 to n.

Definition: A discrete deformation retraction of a voxel set S onto a subset $T \subset S$ is a sequence $S_0, \ldots, S_n$ of voxels sets such that

1. $S_0 = S, S_n = T$ and $T \subset S_i$ for all i

2. $S_0 \approx S_1 \approx \cdots \approx S_n$

In Figure 7 we show an example where our deformation retraction approach guarantees correct topology preservation.
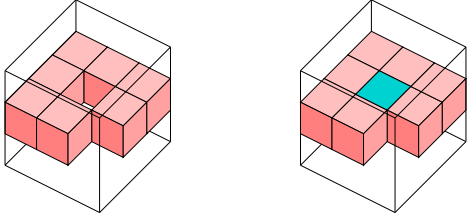
Figure 7. Methods for topology preserving voxel growing described previously attempt to detect topology changes by analyzing the number of connected components of a $3 \times 3 \times 3$ configuration. As shown above, this is not sufficient: In both cases the number of connected components is the same (for the solid as well as for the empty voxels), although the Euler characteristic of the two configurations is different.

3 Morphology

Morphological operators have been in wide use in the image processing community, for a thorough exposition see e.g. [19]. Here we use them to describe distance fields that steer the inflation process of a voxel set.
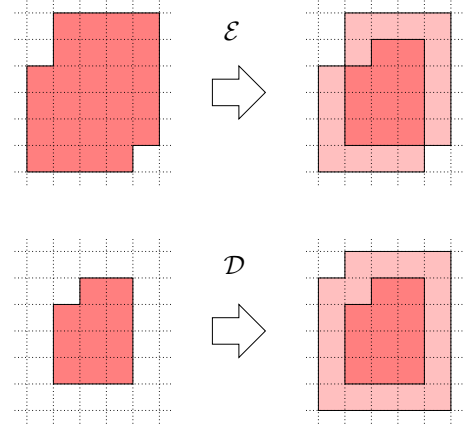
3.1 Erosion and dilation

Intuitively, *erosion* describes the process of removing the boundary layer of a voxel set. *Dilation* on the other hand can be described as pasting another layer of voxels onto a voxel set. To be mathematically more precise let the erosion operator $\mathcal{E}$ and the dilation operator $\mathcal{D}$ be defined as

$$\mathcal{E}(S) = \{ \mathbf{s} \in S : \mathbf{s} \not\sim_e \mathbf{w} \text{ for all } \mathbf{w} \in \mathbb{Z}^3 \}$$
$$\mathcal{D}(S) = \{ \mathbf{w} \in \mathbb{Z}^3 : \mathbf{w} \sim_s \mathbf{s} \text{ for some } \mathbf{s} \in S \}$$

Figure 8 depicts the effect of these operators. Note that $\mathcal{E}$ and $\mathcal{D}$ are in general not inverse to each other,

$$\mathcal{E} \circ \mathcal{D} \neq id \neq \mathcal{E} \circ \mathcal{D}.$$

$\mathcal{E} \circ \mathcal{D}$ and $\mathcal{E} \circ \mathcal{D}$ are also known as the *opening* and *closing* operations.

Figure 8. Morphological operators: Top: The erosion operator $\mathcal{E}$ removes the boundary layer of voxels from a given voxel set. Bottom: The dilation operator $\mathcal{D}$ adds another layer of voxels to a voxel set. As can be seen above, erosion and dilation are not inverses of each other.

Often, the morphological operators are restricted by some other voxel set R. More precisely, we define

$$\mathcal{E}_R(S) = R \cup \mathcal{E}(S)$$

for voxel sets $R \subset S$ and

$$\mathcal{D}_R(S) = R \cap \mathcal{D}(S)$$

for voxel sets $R \supset S$. These restricted operators are illustrated in Figure 9. Repeated erosion/dilation of a voxel set S generates nested sequences of the form

$$S \supset \mathcal{E}_R(S) \supset \mathcal{E}_R^2(S) \supset \cdots$$
$$S \subset \mathcal{D}_R(S) \subset \mathcal{D}_R^2(S) \subset \cdots$$

Following [9] we use these sequences to define two distance functions on the voxel grid. In particular let

$$d_{\mathcal{D},R,S}(\mathbf{s}) = n \quad \text{for} \quad \mathbf{s} \in \mathcal{D}_R^n(S) \setminus \mathcal{D}_R^{n-1}(S)$$

be the distance of the voxel $\mathbf{r} \in R$ from S and let

$$d_{\mathcal{E},R,S}(\mathbf{s}) = n \quad \text{for} \quad \mathbf{s} \in \mathcal{E}_R^n(S) \setminus \mathcal{E}_R^{n-1}(S)$$

be the distance of the voxel $\mathbf{s} \in S$ from the surrounding empty space. Figure 10 depicts these distance functions in the two dimensional setting.

3.2 Topology preserving inflation

Let S be a voxel set and $R \supset S$. Our goal is to describe a discrete topology preserving inflation of S to R. In general,
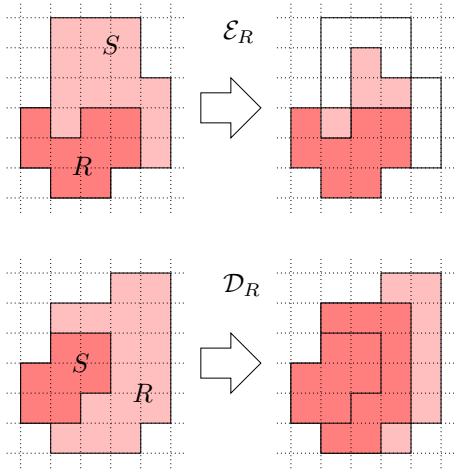
Figure 9. Restricted morphological operators. Top: The restricted erosion operator $\mathcal{E}_R(S)$ removes only those boundary voxels from S that do not belong to R. **Bottom:** The restricted dilation operator $\mathcal{D}_R(S)$ only adds those boundary voxels to S that also belong to R.

Figure 10. Distance functions. Darker colors indicate greater distances. **Left:** Initial configuration. **Middle:** Distance field of $d_{\mathcal{D},R,S}$. **Right:** Distance field of $d_{\mathcal{E},R,S}$

such an inflation does not have to exist because there might be no discrete deformation retraction from R onto S. The best we can do is to find a maximal subset $R' \subset R$ such that there exists a discrete deformation retraction from R' to S. The basic idea is simple: We successively label voxels from R as solid as long as the topology does not change. To be more precise, let us call a voxel $\mathbf{t}$ feasible with respect to a voxel set T, if $T \approx T \cup \{\mathbf{t}\}$. Now we set $R_0 := S$ and let

$$R_{i+1} := R_i \cup \{\mathbf{s}_i\}, \tag{1}$$

where $\mathbf{s}_i \in R \setminus R_i$ is feasible with respect to R_i. We repeat this process until there are no more feasible voxels left and end up with a discrete deformation retraction from $R' := R_n$ to S. Obviously, R' is maximal and we write
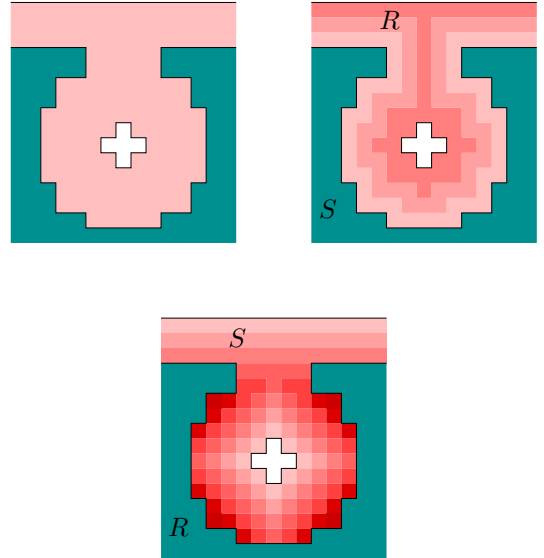
$$S \xrightarrow{R} R'.$$

In general, R' is not uniquely defined by the above procedure. To make the algorithm deterministic we assign a priority $p(\mathbf{s})$ to each voxel $\mathbf{s}$. Then we modify the choice in equation (1) to always select a feasible voxel with maximal priority.

The priority function p can be adapted to different requirements. A particularly intuitive choice is setting

$$p := d_{\mathcal{E},R,S}.$$

The effect of this choice of p is demonstrated in Figure 13. S is a voxel set homeomorphic to a sphere, $\chi(S) = 1$, but R has a handle, $\chi(R) = 0$. The inflation process has to avoid

the handle by placing a cut somewhere. Because p measures the distance of a voxel from the surrounding empty space, voxels which are on the boundary of R are added last during inflation. Hence, the cut is automatically placed at that location of R that has minimal diameter.

Implementation note

Although R' is maximal it is still possible that a rather large part of R is not covered by R'. This can happen, e.g. when the inflation process runs into a dead end as depicted in Figure 11. A straightforward method to overcome this inherent problem is to double the resolution of the voxel sets before applying the inflation process. Similar to Section 2.3/Figure 6 we try to set the eight subvoxels sequentially.
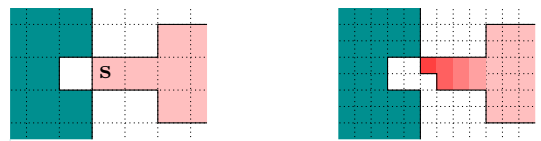
Figure 11. Dead ends may stop the inflation process at early stages. On the left the dark voxels have already been conquered by the inflation process. The center voxel s cannot be conquered, since this would change the topology. Increasing the resolution by subdividing the entire voxel space fixes the problem (right).

6

4 Application Example

In this section we present an application example how to make use of topology preserving inflation in the context of medical image processing.

4.1 Problem statement

Let us consider a discrete intensity-field h defined on $\mathbb{Z}^3$,

$$h : \mathbb{Z}^3 \to \{0, \ldots, m\}.$$

Such data typically arises in medical imaging applications, like e.g. CTs or MRTs. As a typical example we consider a set of MRT scans of a human brain (Figure 16), however, the method is of course applicable to other organs as well.

The intensities of a MRT image give a rough classification of the tissue type of the corresponding pixels. High intensity values correspond to white matter (WM), midrange intensities to grey matter (GM) and low intensities to the surrounding cerebrospinal fluid (CSF). However, it is in general not possible to simply threshold the pixel values in order to exactly classify the pixels: Because of the finite resolution of the imaging device, the pixels actually represent the average intensity taken over a certain volume of the brain.

From what is known from neuro-anatomy we expect the interface of CSF and GM to be homeomorphic to a sphere. However, because of noise in the data and because of cavities that are generated by folds of the cortex, a simple thresholding of the intensity values generally ends up with a surface whose characteristic is different. To get a topologically correct reconstruction of the interface, we proceed in several stages that are explained in the following sections.

4.2 Presegmentation and Classification

In a first step we apply a rather standard presegmentation algorithm to the voxel data. This is done very similar as described in [20]. For the convenience of the reader, we recap the main steps by means of Figure 16.

First the measurement noise which is inherent in every physical data acquisition process is filtered out by applying an anisotropic diffusion filter to the data set. Next the image is segmented by applying an LoG (Laplacian of Gaussian) Operator,

$$M := \{\mathbf{v} \in \mathbb{Z}^3 : LoG(\mathbf{v}) > 0\}.$$

To cut out thin voxel structures, the set M is eroded, and a connected component analysis is performed on $M' := \mathcal{E}(M)$. All but the maximal component are removed and a hole-filling closing operator is applied which leaves us with a mask M'' that approximately covers the GM and WM. By discarding all voxels in the original dataset that are not covered by M'', the skull is automatically removed.

This automatic procedure is an alternative to projecting the MR volume into Talairach space and then apply standard masks to remove the cranium and the cerebellum [22].

The second step is to classify each of the remaining voxels according to the tissue types WM, GM and CSF. Let h be the intensity field and let

$$U_\phi = \{\mathbf{v} \in \mathbb{Z}^3 : h(\mathbf{v}) > \phi\}.$$

We choose two suitable thresholds $\sigma < \tau$ and set

$$
\begin{aligned}
W &:= U_\tau \quad \text{(white matter)} \\
G &:= U_\sigma \setminus W \quad \text{(grey matter)} \\
E &:= \mathbb{Z}^3 \setminus G \setminus W \quad \text{(empty space)}.
\end{aligned}
$$

4.3 Extracting the CSF/GM interface

The previous steps leave us with a set of voxels W, that are classified as white matter. This voxel set provides an accurate initial representation of the brain structure and serves as a starting point for the inflation procedure.

As W is in general not simply connected, we first perform a component analysis and remove all but the maximal component. Inclusions are then removed analogously by applying a component analysis to $\mathbb{Z}^3 \setminus W$. Hence, in the following we can assume, that W is connected and free of inclusions. Still, the topology of W is in general not correct. We remove the handles as described in Section 3.2 by first computing the priority function $p = d_{\mathcal{E}, W, \emptyset}$, then selecting a single voxel $\mathbf{v}$ with maximal priority and inflating this voxel to W,

$$\mathbf{v} \xrightarrow{W} W'.$$

The inflated voxel set W' lies within W and is of the correct topology (Figure 14).

Alternatively, we can take the dual approach and inflate the empty space around W. This will remove the handles by filling them up instead of placing a cut. Kriegeskorte et al. [9] compute both solutions and then decide for each handle separately whether to fill it or whether to cut it based on an estimation of the resulting error.

Finally, the white matter W' is inflated to the grey matter G,

$$W' \xrightarrow{G} W''.$$

Here again we take advantage of a priority function such that the handles of G will be cut where the diameter is smallest (Figure 15).

4.4 Improving the Inflation Process

Note that after the classification the values of the (possibly filtered) intensity field h are not used anymore: the inflation process is controlled by the distance function alone.

This behaviour can sometimes lead to counter intuitive results as is demonstrated in Figure 17, where the cut is not placed at the minimum of the intensity function. We can solve this problem by incorporating the scalar field gradient into our growing procedure without any risk of topology artifacts caused by measuring noise.

The idea is as follows: Instead of inflating $W' \xrightarrow{G} W''$ in one step, we successively inflate

$$W' \xrightarrow{U_{\tau-1}} W^1 \xrightarrow{U_{\tau-2}} \ldots \xrightarrow{U_{\sigma+1}} W^{\tau-\sigma} \xrightarrow{G} W'',$$

i.e. we successively lower the σ threshold. Figure 17 demonstrates this process. Using this approach we assure that the topological cuts always are done along the valleys of the intensity field.

4.5 Extracting ∂S

For visualization purposes as well as for other downstream applications it is often necessary to explicitly approximate ∂S, e.g. as a triangle mesh $\mathcal{T}$. Obviously $\mathcal{T}$ should not only approximate ∂S well but it should also be of the same topology. To construct $\mathcal{T}$ we follow the approach of Lachaud [10] and use a marching cubes algorithm [13] whose lookup table has been modified such that it matches the neighborhood relation $\sim_s$ of the solid voxels. These modifications are straightforward, an example is depicted in Figure 12. Note that since $\sim_s$ and $\sim_e$ are compatible, there are no ambiguate voxel configurations.
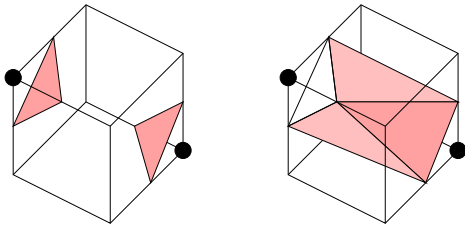
Figure 12. Left: Standard marching cubes triangulation of a voxel with two diagonally opposing marks. Right: Modified triangulation that respects the 26-neighborhood relation.

4.6 Results

We applied our inflation procedure to different MRT scans. Figure 14 shows the reconstruction of the white matter surface from a MRT volume of $256 \times 256 \times 130$ voxels. Figure 15 shows the grey matter surface of the same dataset. The whole resconstruction process took less than 5 minutes on a standard PC.

5 Conclusion and Future work

In this paper we presented a mixed topological / morphological framework to accurately model the inflation process of surfaces in the discrete setting of voxel spaces. We showed the applicability of this framework by embedding it into an algorithm to extract the cortical surface of a brain from a set of MRT scans. By letting the intensity gradient control the inflation process, we assure that separating cuts are placed at places where the intensity field takes on a minimum.
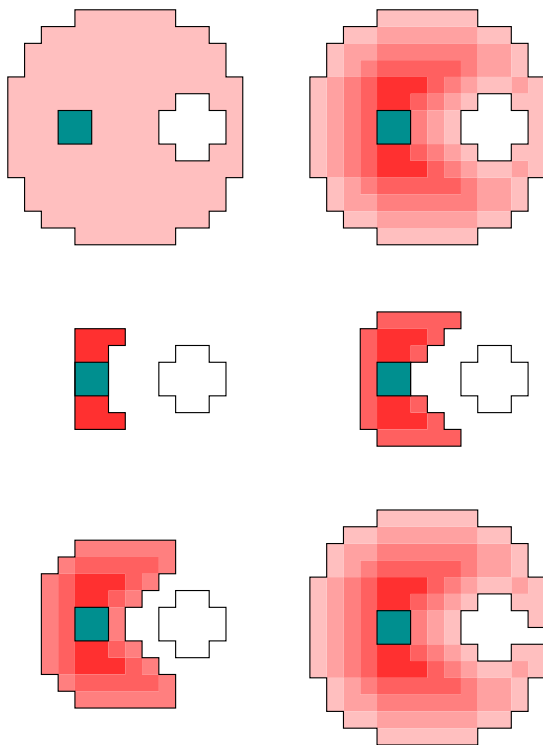
As future work we plan to subject our reconstruction results to a rigorous validation process, as it is commonplace for medical applications. We also plan to apply the reconstruction framework in other application, like the removal of topological noise from volume data that, e.g., was gathered by 'space carving' methods from range images.
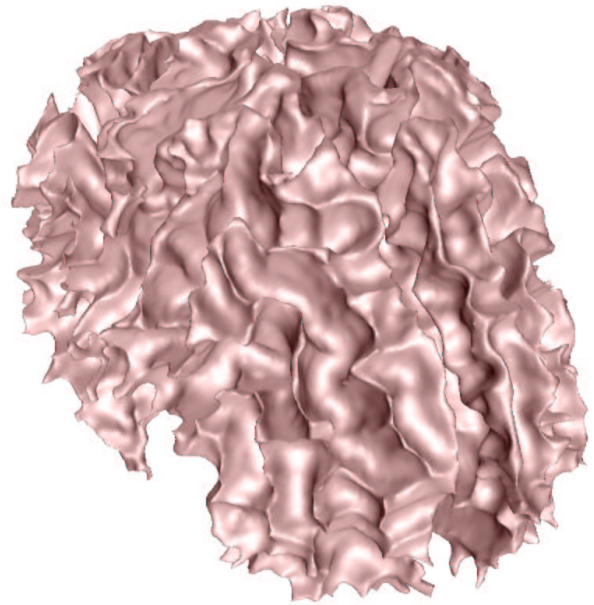
References

[1] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In *SIGGRAPH 96 proceedings*, pages 303–312, 1996.

[2] C. Davatzikos and R. N. Bryan. Using a deformable surface model to obtain a shape representation of the cortex. *IEEE Trans. Med. Imag.*, 15:785–795, 1996.

[3] B. Fischl, A. Liu, and A. M. Dale. Automated manifold surgery: Constructing geometrically accurate and topologically correct models of the human cerebral cortex. *IEEE Trans. Med. Imag.*, 20:70–80, 2001.

[4] G.Bertrand A boolean characterization of three-dimensional simple points. *Pattern Recognition Letters*, 17:115–124, 1996.

[5] I. Guskov and Z. Wood. Topological noise removal. In *GI 2001 proceedings*, pages 19–26, 2001.

[6] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2001.

[7] K. H. Höhne and W. H. Hanson. Interactive 3d segmentation of mri and ct volumes using morphological operations. *J. Comput. Assist. Tomogr.*, 16(2):285–294, 1992.

[8] T. Y. Kong and A. Rosenfeld. Digital topology: Introduction and survey. *Computer Vision, Graphics and Image Processing*, 48:357–393, 1989.

[9] N. Kriegeskorte and R. Goebel. An efficient algorithm for topologically correct segmentation of the cortical sheet in anatomical mr volumes. *NeuroImage*, 14:329–346, 2001.

[10] J.-O. Lachaud. Continuous analogs of digital boundaries: A topological approach to iso-surfaces. *Graphical Models*, 62:129–164, 2000.

[11] L. Latecki and C. C. Ma. An algorithm for a 3d simplicity test. *Computer Vision and Image Understanding*, 63(2):388–393, 1996.

[12] T. C. Lee, R. L. Kashyap, and G. N. Chu. Building skeleton models via 3d medial surface/axis thinning algorithms. *Graphical Models and Image Processing*, 56(4):462–478, 1994.

[13] W. E. Lorensen and H. E. Cline. Marching cubes: a high resolution 3d surface reconstruction algorithm. In *SIGGRAPH 87 proceedings*, pages 163–169, 1987.

[14] D. MacDonald, N. Kabani, D. Avis, and A. C. Evans. Automated 3-d extraction of inner and outer surfaces of cerebral cortex from mri. *NeuroImage*, 12:340–356, 2000.

[15] P.Alexandroff and H.Hopf. *Topologie, Erster Band: Grundbegriffe der mengentheoretischen Topologie*. Springer, 1935.

[16] L. Robert and G. Malandain. Fast binary image processing using binary decision diagrams. *CVIU*, 72(1):1–9, 1998.

[17] A. Rosenfeld. Digital topology. *American Mathematical Monthly*, 86:621—630, 1979.

[18] P. K. Saha, B. B. Chauduri, B. Chanda, and D. Majumber. Topology preservation in 3d digital space. *Pattern Recognition*, 27(2):295–300, 1994.

[19] J. Serra. *Image analysis and mathematical morphology*. Academic Press, 1982.

[20] D. W. Shattuck, S. R. Sandor-Leahy, K. A. Schaper, D. A. Rottenberg, and R. M. Leahy. Magnetic resonance image tissue classification using a partial volume model. *NeuroImage*, 13:856–876, 2001.

[21] R. Stokking, K. L. Vincken, and M. A. Viergever. Automatic morphology-based brain segmentation (mbrase) from mri-t1 data. *NeuroImage*, 12:726–738, 2000.

[22] J. Talairach and P. Tournoux. *Co-Planar Stereotaxic Atlas of the Human Brain*. Thieme Medical Publishers, NY, 1988.

Figure 14. White matter surface: The interface between the white and the grey matter provides the starting point for the inflation process.

Figure 13. Inflation of a voxel set: a) The initial voxel sets S (cyan) and R (pink). b) Visualization of the priority function, darker colors signify greater distance from the boundary. c) - f) Succesive stages of the inflation process. The last voxel cannot be set, since it would cause a topology change.
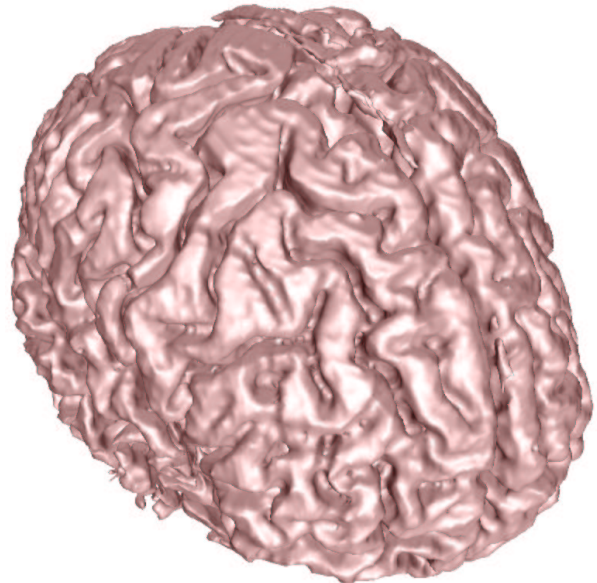
Figure 15. Grey matter surface: By inflating the white matter, we obtain a topologically correct approximation of the interface between grey matter and cerebrospinal fluid.
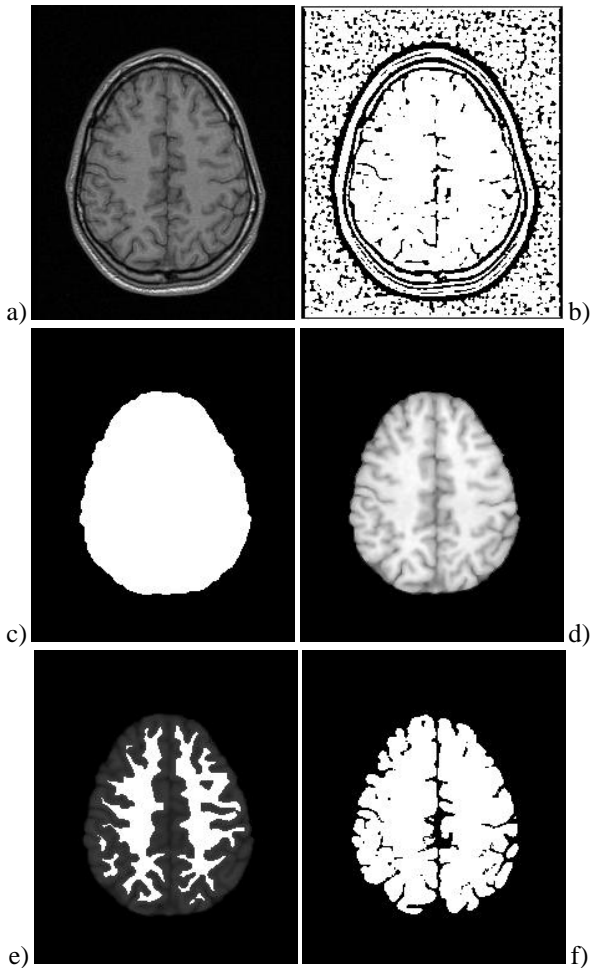
Figure 16. MRT processing pipeline: a) Original data. b) LoG filtered data. c) After morphological segmentation. d) Skull removed. e) Inital white matter approximation. f) After the inflation procedure.
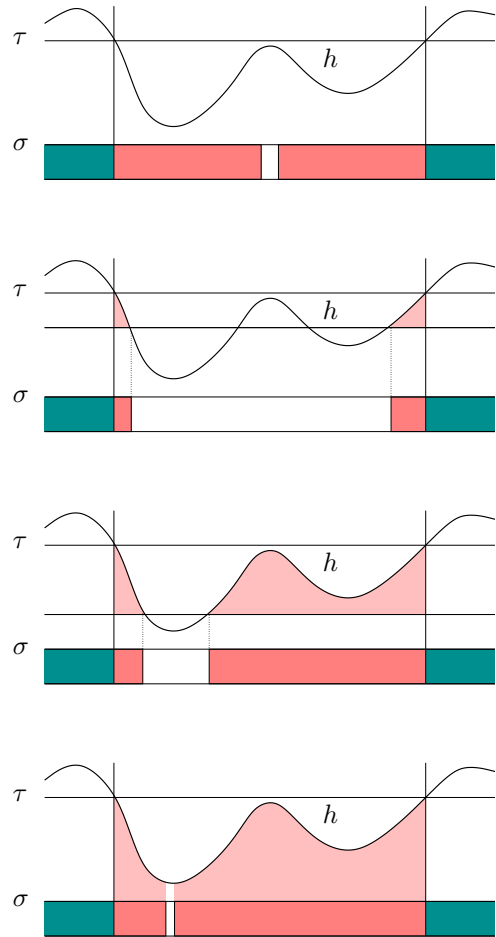

Figure 17. Depicted above is a 1-dimensional slice through a voxel set. If $h(\mathbf{v}) \geq \tau$ the corresponding voxel $\mathbf{v}$ belongs to the white matter W. If $\sigma \leq h(\mathbf{v}) < \tau$, $\mathbf{v}$ belongs to the grey matter G. Using a distance driven inflation process causes the cut to be placed halfway between the boundaries (top). If instead, σ is successively lowered as shown in the lower pictures the inflation process will locate the cut at the minimum of the intensity function.