

Variational Segmentation of Elongated Volumetric Structures*

Christian Reinbacher, Thomas Pock, Christian Bauer, Horst Bischof
Institute for Computer Graphics and Vision
Graz University of Technology, Austria
{reinbacher,pock,cbauer,bischof}@icg.tugraz.at

Abstract

We present an interactive approach for segmenting thin volumetric structures. The proposed segmentation model is based on an anisotropic weighted Total Variation energy with a global volumetric constraint and is minimized using an efficient numerical approach and a convex relaxation. The algorithm is globally optimal w.r.t. the relaxed problem for any volumetric constraint. The binary solution of the relaxed problem equals the globally optimal solution of the original problem. Implemented on today's user-programmable graphics cards, it allows real-time user interaction. The method is applied to and evaluated on the task of articular cartilage segmentation of human knee joints and segmentation of tubular structures like liver vessels and airway trees.

1. Introduction

Medical volumetric images are three-dimensional images which contain several anatomical structures. Depending on the imaging method, different structures can be visualized. These images are usually inspected by radiologists. For diagnosis and therapeutic interventional planning, the first step usually is an accurate segmentation of one object. In many applications (e.g. cartilage) the segmentation of very thin and elongated objects is necessary which poses additional difficulties. Current research is focusing on speeding up this lengthy process by providing either fully-automatic or semi-automatic segmentation methods. Automatic methods like [12, 13, 14, 15] have proven very useful for medical studies where a large amount of data has to be processed. However, in some cases the ability to control and correct segmentation results is necessary in order to gain the required accuracy [19]. In this case the goal is to minimize the required interaction time.

¹This work was supported by the Austrian Science Fund (FWF) under the doctoral program Confluence of Vision and Graphics W1209 and the Austrian Research Promotion Agency (FFG) under the project SILHOUTTE (825843).

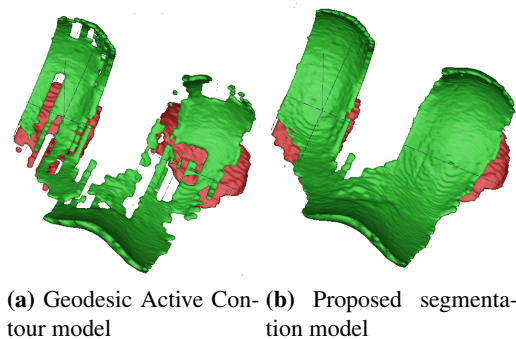


Figure 1: Segmentation results for articular cartilage of a human knee joint. (a) is obtained using a standard GAC segmentation model, (b) is obtained by incorporating edge direction information into the segmentation model.

Modern imaging technologies such as Magnetic Resonance Tomography (MRT) or Computer Tomography (CT) provide a basis for *in vivo* investigations. They deliver a full 3D dataset of a joint or organ, visualizing the main anatomical structures. MRT images however suffer from intensity inhomogeneities caused by inhomogeneities in the static magnetic field and different thickness of tissue. The impact of these inhomogeneities can be reduced by using segmentation methods based on image gradients.

The main contribution of this paper is the use of an anisotropic weighted Total Variation energy with an additional global volume constraint to segment thin and elongated structures like articular cartilage directly in 3D. The volume constraint allows us to define a minimum size of the resulting segmentation. The segmentation model works interactively, allowing the user to incorporate prior knowledge into the segmentation process and correct the segmentation results. The model is solved in a globally optimal manner. The remainder of the paper is organized as follows: Section 2 discusses related work to variational image segmentation. In Section 3 we introduce our segmentation framework and a minimization technique to solve the proposed

segmentation functional. In Section 4 experimental results are given. Finally, we draw conclusions in Section 5.

2. Related Work

In the following, we present some of the methods which are closely related to our proposed variational segmentation framework. One of the most influential works in image segmentation was the introduction of the Snake model by Kass *et al.* [16]. A contour is deformed by internal and external forces. A major drawback of this approach is the local nature of the evolution process which is liable to get stuck in local minima. Also the topology of the contour can not be changed easily. To overcome these limitations Caselles *et al.* [6] proposed the so called Geodesic Active Contour (GAC) model which aims to minimize the following energy

$$\inf_C \left\{ E_{GAC}(C) = \int_0^{|C|} g(|\nabla I(C(s))|) ds \right\}, \quad (1)$$

where $|C|$ is the Euclidean length of the curve C and $g(x)$ an edge detection function which is close to 0 at a strong edge in the image I . A frequent choice is $g(|\nabla I|) = e^{-\beta|\nabla I|^\alpha}$, for some reasonable values of α and β . In a discrete setting, the energy in (1) is similar to the energy proposed by Boykov and Kolmogorov [3], which they minimized using Graph Cuts. The GAC model integrates the Euclidean length of a curve with an energy term derived from image boundaries. In contrast to the Snakes model this contour is allowed to change its topology during evolution. Nevertheless the minimization based on Level-Sets is prone to stuck in local minima. In a continuous formulation, the GAC energy can be realized with the weighted Total Variation (TV) as shown by Bresson [4]. He proposed to minimize the following energy

$$\inf_{u \in \{0,1\}} \left\{ \int_{\Omega} g(x) |\nabla u| dx \right\}, \quad (2)$$

with the image domain Ω , and $u : \Omega \rightarrow \{0, 1\}$ the segmentation. Bresson showed that by letting u vary continuously in the interval $[0, 1]$ (2) becomes a convex functional which can be solved globally optimal. This makes the resulting segmentation independent from the initialization. The *TVSeg* approach of Unger *et al.* [22] adds the possibility to incorporate user information into the segmentation process by minimizing

$$\inf_{u \in [0,1]} \left\{ \int_{\Omega} g(x) |\nabla u| dx + \int_{\Omega} \lambda(x) |u - f| dx \right\}, \quad (3)$$

where $f : \Omega \rightarrow \mathbb{R}$ is a user provided potential function and $\lambda(x)$ a spatially varying parameter which determines the balance of the length term and the data term.




Figure 2: Incorporating edge direction into the weighted Total Variation term. The relation between primal and dual variable are depicted for the isotropic segmentation in (a), and for the anisotropic segmentation model in (b)

3. The Segmentation Framework

Although the segmentation methods described in Section 2 are very powerful and applicable to a variety of tasks, they are not well suited for the task of segmenting elongated objects. The weighted regularization term minimizes the surface of the resulting segmentation while neglecting edge direction information and thus prefers objects with a low surface to volume ratio. Furthermore, in many segmentation tasks, the size of the objects of interest is known. This information could be used as additional information in the segmentation process. In order to circumvent these problems we propose to solve the following minimization problem:

$$\inf_{u \in K^0} \left\{ \int_{\Omega} \sqrt{\nabla u^T D(x) \nabla u} dx + \lambda \int_{\Omega} u f dx \right\}, \quad (4)$$

where

$$K^0 = \left\{ u : \Omega \rightarrow \{0, 1\} : \int_{\Omega} u dx \geq c \right\}, \quad (5)$$

and $c : 0 \leq c \leq \int_{\Omega} dx$ is a volumetric constraint which defines a minimum size of the segmentation.

3.1. Length Term

The length term is an anisotropic Total Variation term which was also used recently by Olsson *et al.* [20]. $D(x)$ is a symmetric non-singular square matrix which specifies the local metric at a given point x . Setting $D(x)$ to the identity matrix results in the regularization being the Total Variation of u and results in minimizing the surface area. By setting

$$D(x) = \text{diag} (g(x)^2), \quad (6)$$

(4) becomes a weighted Total Variation. $I : \Omega \rightarrow \mathbb{R}$ is the image we want to segment. The edge map also contains information about the direction of the edges, which we completely neglected so far. To add this information, we define

$D(x)$ for our segmentation framework as

$$D(x) = g(x)nn^T + n_0n_0^T + n_1n_1^T, \quad (7)$$

where $n = \frac{\nabla I}{\|\nabla I\|}$, n_0 is the tangent of n and $n_1 = n \times n_0$.

The difference between $D(x)$ defined in (6) and (7) is depicted in Fig. 2. The red line represents an edge. $D(x)$ as defined in (6) describes a circle which is scaled with $g(x)$. We refer to this segmentation model as *Isotropic Model* in the following. $D(x)$ as defined in (7) describes an ellipse with its major axis oriented in the direction of the local edge. The minor axis is scaled with $g(x)$. We refer to this segmentation model as *Anisotropic Model* throughout the paper. The effect of this modification is that ∇u is more likely to be oriented in the edge direction during the minimization process.

3.2. Data Term

The potential function f can be used to define an affinity of the image belonging to the foreground ($u \rightarrow 1$) or to the background ($u \rightarrow 0$) region. Setting $f < 0$ results in u tending towards positive values during the minimization process. $f > 0$ results in u tending towards negative values. $f = 0$ means that no information about fore- or background is given in this area and the pure geodesic energy is minimized.

As seen in the review about level-set methods for image segmentation by Cremers [11] the potential function f is usually given in a form of

$$f = -\log \frac{P(B)}{P(F)}, \quad (8)$$

where $P(F)$ denotes the probability for a region to belong to the foreground and $P(B)$ the probability for a region to belong to background. For example, setting $P(F) = e^{-\frac{(I-\mu_F)^2}{2\sigma^2}}$ and $P(B) = e^{-\frac{(I-\mu_B)^2}{2\sigma^2}}$ with μ_F and μ_B being the mean gray values of foreground and background regions one arrives at

$$f = -\log P(F) + \log P(B) \quad (9)$$

$$= \frac{1}{2\sigma^2} ((I - \mu_F)^2 - (I - \mu_B)^2), \quad (10)$$

which equals the data-driven term used by Chan and Vese [10] for their Mumford-Shah like segmentation functional, when neglecting the constant factor $\frac{1}{2\sigma^2}$. In our segmentation framework we define $P(F)$ and $P(B)$ by Parzen windowing gray value histograms of user-defined foreground and background regions.

3.3. Solving the Segmentation Model

In this section we present numerical algorithms to solve the minimization problem (4).

Discrete Setting

For the implementation in a computer, functions and operators have to be discretized. The image domain Ω can be discretized by a regular pixel grid of size $(X \times Y \times Z)$

$$\Omega^h = \{(ih_x, jh_y, kh_z) : 0 \leq i < X, 0 \leq j < Y, 0 \leq k < Z\}, \quad (11)$$

where h_x, h_y, h_z denote the discretization widths. The positions in the grid are defined as (i, j, k) . The discrete version of the three-dimensional gradient operator is approximated using simple forward difference with Neumann boundary conditions. From now on we will restrict our attention to the discrete setting.

Minimization

It is well known that functionals like (4) are hard to minimize due to the L^1 norm $|\nabla u|$. Carter [5], Chan *et al.* [9], Chambolle [7] propose a dual formulation of the Total Variation. First, we can rewrite the weighted Total Variation term as

$$\sqrt{\nabla u(i, j, k)^T D(i, j, k) \nabla u(i, j, k)} = \left| D(i, j, k)^{1/2} \nabla u(i, j, k) \right|. \quad (12)$$

Then, by applying the Legendre-Fenchel transformation, (12) can be written as

$$\left| D(i, j, k)^{1/2} \nabla u(i, j, k) \right| = \max_{\|p(i, j, k)\| \leq 1} \left\langle p(i, j, k), D(i, j, k)^{1/2} \nabla u(i, j, k) \right\rangle, \quad (13)$$

which leads to the following saddle point formulation of our image segmentation energy

$$\min_{u \in K^0} \max_{p \in C} \left\{ \left\langle p, D^{1/2} \nabla u \right\rangle + \langle u, \lambda f \right\}, \quad (14)$$

with the convex set

$$C = \{p : \Omega \rightarrow \mathbb{R}^3 \mid \|p(i, j, k)\| \leq 1, \forall i, j, k \in \Omega\}. \quad (15)$$

Convex Relaxation

The saddle point problem in (14) is convex, but carried out over a non-convex set K^0 . A common way to make sets like K^0 convex is by letting u vary continuously in the interval $[0, 1]$. This results in

$$K = \{u : \Omega \rightarrow [0, 1] : Su \geq c\}, \quad (16)$$

with S being a $1 \times (XYZ)$ vector of ones. Now we can apply the convergent primal-dual algorithm of Pock *et al.*

[21] and arrive at the following update rules for u and p :

$$\begin{cases} p^{n+1} = \Pi_C \{p^n + \tau_D D^{1/2} \nabla \bar{u}^n\} \\ u^{n+1} = \Pi_K \{u^n - \tau_P (-\operatorname{div}(D^{1/2} p^{n+1}) + \lambda f)\} \\ \bar{u}^{n+1} = 2u^{n+1} - u^n, \end{cases} \quad (17)$$

with $\bar{u}^0 = u^0 = 0$, $p^0 = 0$ and τ_D and τ_P being the time steps for the dual and primal update respectively. According to [21] the time steps are chosen so that $\tau_D \tau_P \leq 1/L^2$ with $L^2 = \|\operatorname{div}\|_2^2 = 4(1/h_x^2 + 1/h_y^2 + 1/h_z^2)$ for a three-dimensional input image.

Computing the projections

The reprojection of the dual variable on the convex set C can be achieved by a pointwise orthogonal projection of the dual variable onto the unit ball

$$\Pi_C p(i, j, k) = \frac{p(i, j, k)}{\max\{1, \|p(i, j, k)\|\}}. \quad (18)$$

The reprojection of the primal variable onto the convex set K needs more attention since we have to take non-local constraints into account. Basically, the projection can be written as computing the minimizer of

$$\min_v \left\{ \frac{1}{2} \|u - v\|^2 \right\} \text{ s.t. } v \in [0, 1], Sv - c \geq 0, \quad (19)$$

If the second constraint is omitted, i.e. $c = 0$ the solution of (19) is given by a simple pointwise clamping of v to the interval $[0, 1]$. In order to solve the full projection we introduce the Lagrange multiplier η to account for the volume constraint and reformulate the constrained minimization problem (19) as the following saddle point problem:

$$\min_{v \in [0,1]} \max_{\eta \geq 0} \left\{ \langle \eta, Sv - c \rangle + \frac{1}{2} \|u - v\|^2 \right\}. \quad (20)$$

This problem can again be solved using the primal dual algorithm of [21].

$$\begin{cases} \eta^{n+1} = \eta^n + \sigma_D (S\bar{v} - c) \\ v^{n+1} = \max \left\{ 0, \min \left\{ 1, \frac{v^n + \sigma_P (\eta^{n+1} + u/\tau_P)}{1 + \sigma_P/\tau_P} \right\} \right\} \\ \bar{v}^{n+1} = 2v^{n+1} - v^n, \end{cases} \quad (21)$$

with $\bar{v}^0 = v^0 = u$, $\eta^0 = 0$ and the timesteps $\sigma_P \sigma_D \leq 1/|S|$. This algorithm is iterated until η falls below a certain threshold ϵ .

Computing a Binary Segmentation

A binary segmentation is obtained by thresholding u with $t \in [0, 1]$. Chan *et al.* [8] showed, that a thresholded u is

a global minimizer of the original problem for almost every t . For our segmentation model this holds only when c is set to zero. Lempitsky *et al.* [17] propose a rounding algorithm called pinpointing which verifiably performs better than simple thresholding to obtain an integer solution. However, this algorithm has too high computational complexity for our interactive 3D segmentation task. Thus, we present the user different thresholds of u simultaneously. In Fig. 3 an example of this thresholding is depicted. In the upper part of the bone the different thresholds are far apart, whereas in regions with a distinct segmentation border these thresholds close to each other. This output can be seen as additional cue for the user which tells him in which areas further input is needed in order to obtain a binary segmentation.




Figure 3: Non-binary segmentation: We present the user a differently thresholded u . In the upper part of (a) the segmentations (depicted in red) are far apart. In (b) the corresponding u is shown.

4. Results

In the following section we will present experimental results of our segmentation framework applied to real-world datasets. We will focus on accuracy as well as speed compared to manual segmentations.

All experiments were carried out on a single PC with an Intel Core 2 Quad processor at 3,4 GHz and 4 GB RAM, equipped with a NVidia GeForce Tesla graphics card with 4 GB video memory, running a recent 64 bit Linux distribution. Utilizing this graphics card we are able to handle volumes of $512 \times 512 \times 300$ voxels, which are quite large datasets. We are able to achieve user interactivity with frame rates up to 20 fps. We implemented our segmentation framework using a recent version of CUDA (2.3). A GUI was created, which allows the user to load volumes in file formats common for medical imaging (e.g. DICOM, Analyze), navigate through the dataset, draw and erase weak (f like in (8)) and hard ($f = \infty$ and $f = -\infty$ resp.) foreground and background regions, save and load segmentations for later analysis and save and load constraints. All 3D renderings were created using ITK-SNAP [23].

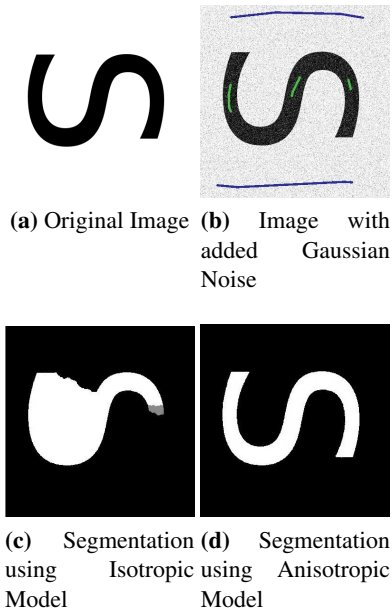


Figure 4: Comparison of the segmentation models applied to a synthetic image. In (b) foreground regions are depicted in green, background regions in blue. These regions were provided to both segmentation models. The proposed segmentation model is able to segment the distorted image correctly.

4.1. Synthetic Example

In Fig. 4 we use an synthetic image to show the influence of including anisotropic edge information in the segmentation process. In Fig. 4(a) the original image is depicted, Fig. 4(b) shows the image distorted with Gaussian noise, Fig. 4(c) the segmentation with the GAC model and Fig. 4(d) the segmentation with the proposed model. As one can see, the GAC prefers convex regions with a minimal length periphery, whereas the additional directional information also allows to segment the whole “S” correctly.

4.2. Cartilage Segmentation

Segmentation of articular cartilage is a hard task for various reasons. In addition to the imaging artifacts, cartilage is a very thin material and its contrast to surrounding soft tissue like ligaments or muscles is low. Especially the ligaments between the medial and lateral condyles of the femur have the same gray value range as cartilage.

We evaluated our segmentation method using MRT datasets of human knee joints. All datasets were acquired using a 3T scanner with a T_2 -weighted Ds3D sequence. The voxel resolution is $0.29 \times 0.29 \times 0.6\text{mm}^3$. The size of the datasets is $448 \times 512 \times 160$ voxels. We set the timesteps τ_D and τ_P for the minimization of the segmentation model

(4) to 0.5. The parameters $\alpha = 15.0$ and $\beta = 0.55$ for the edge detection function yield the best results for the bone segmentation for all datasets. The weighting factor λ was set to 0.005. Although we are able to process the whole dataset in terms of memory consumption, it is advisable to crop the volume to include only the relevant data, to speed up the segmentation process.

The gray value range of the datasets is scaled to the interval $[0, 1]$. However, the user may alter the window level and size to maximize the contrast of the desired object. The user defines foreground and background regions by simple drawing on a slice of the dataset with the mouse. As the user provides more input, the variable u will evolve to the optimal solution. This evolution process can be observed, while the algorithm remains interactive. After the first input is provided, it usually takes several seconds until the optimal solution is found. When modifying the segmentation, the algorithm reacts very fast, even for large datasets.

Directional Edge Information

In order to show the impact of using directional edge information, we segmented the cartilage of a whole joint in 3D with directional edge information. We then switched to the GAC model, leaving the user-provided foreground and background regions constant. In Fig. 1 the differences between the two models is depicted using 3D renderings from the same viewpoint. As one can see, the addition of directional edge information requires less user interaction to segment objects with a high surface to volume ratio, like articular cartilage, successfully. Figure 5 shows the segmentation result of another data set. All segmentation models in Fig. 5(b), (c) and (d) were provided the same user input in (a). Only the anisotropic segmentation model is able to segment the cartilage surfaces with a reasonable amount of user input.

Speed Comparisons

For speed examinations, we segmented several datasets and aimed for an accurate segmentation of the full cartilage of femur and tibia, with special focus on the medically relevant contact area of these surfaces. We measured the time until we were satisfied with the segmentation result for each cartilage and bone separately.

We can conclude that an experienced user is able to segment the femoral and tibial bones and cartilage surfaces of a human knee joint in 30-40 minutes. A manual segmentation using simple drawing tools and specialized touch displays usually lasts 2-3 hours, as shown by McWalter *et al.* [18]. They were able to segment the cartilage of patella and tibia in 60 minutes each. Note that the segmentation of the femoral cartilage is difficult because of the ligaments between the medial and lateral condyles. We estimate an

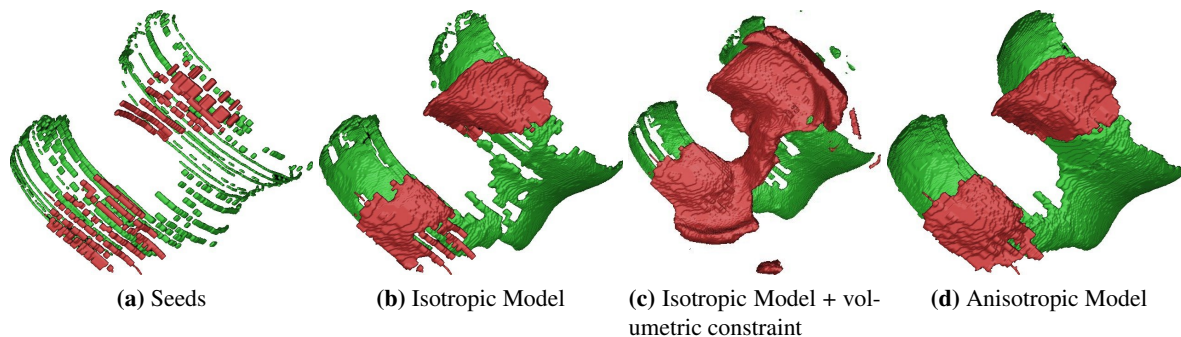


Figure 5: Segmentation results of the cartilage surfaces of a human knee joint. The user input in (a) was provided to both segmentation models. While this is sufficient for the anisotropic segmentation model to correctly segment the cartilage (d), the Isotropic Model fails (b). An additional volume constraint causes the Isotropic Model to leak out (c).

other hour for the segmentation of this cartilage surface using their method.

Using the Volume Constraint

The global volumetric constraint enables the user to define a minimum size of the segmentation. The user interface of the designed segmentation application shows the user the current segmentation size. Within a study, the volume of bone and cartilage stays roughly the same. Thus, after the first successful segmentation the user sets c to the current segmentation size for the next segmentation. This reduces the amount of necessary user input. In our Experiments we found out that 10-15% of the voxels belonging to femoral cartilage have to be provided by the user, and 15-20% of the voxels belonging to tibial cartilage. With a volumetric constraints these numbers reduced by 5%. Note that in Fig. 5(c) the volume constraint causes the Isotropic Model to leak out.

4.3. Tubular Object Segmentation

We applied our segmentation framework to the task of segmenting medical tubular objects like the airway tree of the human lung and blood vessels of the liver. The experiments were carried out using CT datasets with an isotropic voxel size of 0.57mm and an image size of $512 \times 512 \times 300$. To assist the user in providing foreground and background regions we applied a tubular structure detector based on an approach proposed by Bauer *et al.* [1, 2]. The detections were used as foreground regions, background regions were created by thresholding a distance transformation of the detections. Figure 7(a) shows the initial seeds for the segmentation of a liver vessel dataset. Figure 7(b) shows the segmentation result using the isotropic segmentation model, (c) with an additional volume constraint and (d) the segmentation with the Anisotropic Model. As one can clearly see in

the magnification, the Anisotropic Model is able to segment the fine details of the vessels near the border of the liver.

The segmentation of one human airway is shown in Fig. 6. Figure 6(b) shows the segmentation result of the GAC model. In Fig. 6(c) an additional volume constraint c was provided to force the segmentation to reach also the fine structures at the end of the airway. The Anisotropic Model is able to segment also the fine details in the bottom right of Fig. 6(d).

5. Conclusion

This paper presented a interactive segmentation framework for thin volumetric structures. This framework combines local anisotropic edge information with a global volume constraint. In contrast to the energy formulation of Boykov and Kolmogorov [3] we are able to formulate a non-local volumetric constraint c which can be very useful for a variety of segmentation tasks where the size of the object is roughly known. Due to the continuous formulation we do not suffer from metrication errors during the minimization process. The used minimization algorithm is proven to find the globally optimal solution for our segmentation problem. The iterative minimization of the segmentation model can easily be computed in parallel because in each update step every voxel only depends on its neighboring voxels. This perfectly suits the implementation on a modern user-programmable graphics card.

In Section 4 we applied our segmentation method to real-world datasets. We were able to achieve a considerable speedup compared to manual segmentation methods. In our experiments we found out that the limiting factor was the user interface. Future work will have to deal with improving the user interaction with three-dimensional data sets and visualization of the intermediate segmentation result directly in 3D. As shown in Section 4.3 the potential function may not only be provided by the user but also by specifically de-

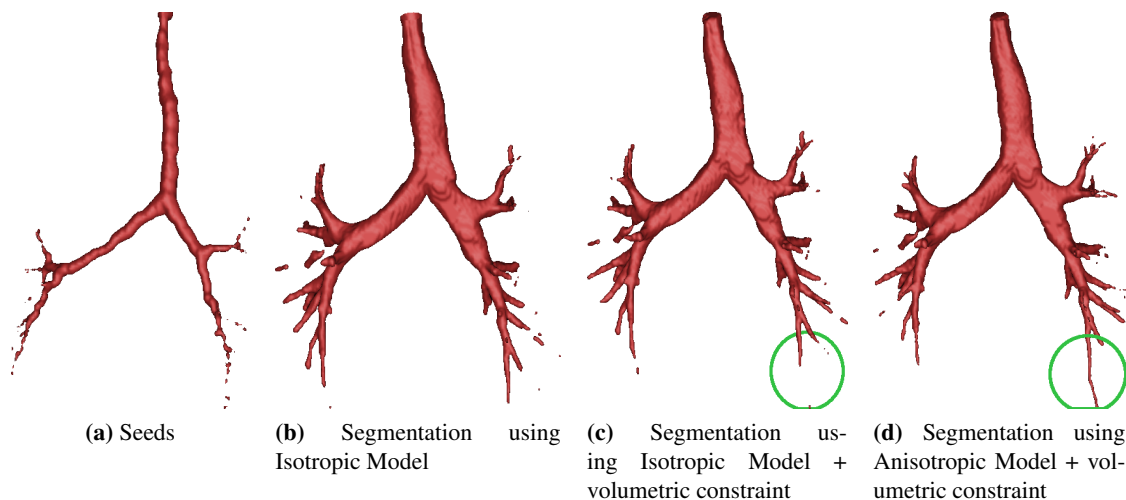


Figure 6: Segmentation of an airway tree: Comparison of different segmentation models. Note the fine details which can only be segmented using directional edge information in the segmentation model.

signed detection algorithms. A statistical model based on previous segmentations which improves with every newly segmented object would also help to minimize user interaction and will be part of our future research.

References

- [1] C. Bauer, H. Bischof, and R. Beichel. Segmentation of airways based on gradient vector flow. In *International Workshop on Pulmonary Image Analysis*, Medical Image Computing and Computer Assisted Intervention, pages 191–201, sep 2009. 6
- [2] C. Bauer, T. Pock, E. Sorantin, H. Bischof, and R. Beichel. Segmentation of interwoven 3d tubular tree structures utilizing shape priors and graph cuts. *Medical Image Analysis*, 14:174–184, 2010. 6
- [3] Y. Boykov and V. Kolmogorov. Computing geodesics and minimal surfaces via graph cuts. In *Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on*, volume 1, pages 26–33, 2003. 2, 6
- [4] X. Bresson. *Image segmentation with variational active contours*. PhD thesis, EPFL, Lausanne, 2005. 2
- [5] J. Carter. *Dual Methods for Total Variation-based Image Restoration*. PhD thesis, UCLA, Los Angeles, 2001. 3
- [6] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *International Journal for Computer Vision*, 77:423–451, 1997. 2
- [7] A. Chambolle. An algorithm for total variation minimizations and applications. *Journal of Mathematical Imaging and Vision*, 10(1-2):89–97, 2004. 3
- [8] S. Chan, T. F. Esedoglu and M. Nikolova. Algorithms for finding global minimizers of image segmentation and denoising models. *SIAM Journal on Applied Mathematics*, 66:1632–1648, 2006. 4
- [9] T. Chan, G. Golub, and P. Mulet. A nonlinear primal-dual method for total variation-based image restoration. *SIAM Journal of Applied Mathematics*, 20(6):1964–1977, 1999. 3
- [10] T. Chan and L. Vese. Active contours without edges. *Image Processing, IEEE Transactions on*, 10:266–277, 2001. 3
- [11] D. Cremers, M. Rousson, and R. Deriche. A review of statistical approaches to level set segmentation: integrating color, texture, motion and shape. *International Journal of Computer Vision*, 72:195–215, 2007. 3
- [12] E. Dam, J. Folkesson, M. Loog, P. Pettersen, and C. Christiansen. Efficient automatic cartilage segmentation. In *Joint Disease Workshop*. ITU, 2006. 1
- [13] J. Folkesson, E. B. Dam, O. F. Olsen, P. C. Pettersen, and C. Christiansen. Segmenting articular cartilage automatically using a voxel classification approach. *IEEE Transactions on Medical Imaging*, 26(1):106–115, Jan. 2007. 1
- [14] J. Fripp, S. Crozier, S. Warfield, and S. Ourselin. Automatic segmentation of articular cartilage in magnetic resonance images of the knee. In *10th International Conference on Medical Image Computing and Computer Assisted Intervention*, pages 186–194, Brisbane, Australia, 2007. 1
- [15] J. Kaftan, H. Tek, and T. Aach. A two-stage approach for fully automatic segmentation of venous vascular structures in liver ct images. In J. P. W. Pluim and B. M. Dawant, editors, *Medical Imaging 2009: Image Processing*, page 725911, 2009. 1
- [16] M. Kass, A. Witkins, and D. Terzopoulos. Snakes: Active contour models. *International Journal for Computer Vision*, 1(4):321–331, January 1987. 2
- [17] V. Lempitsky, P. Kohli, C. Rother, and T. Sharp. Image segmentation with a bounding box prior. In *IEEE International Conference on Computer Vision*, 2009. 4
- [18] E. McWalter, W. Wirth, S. M., R. von Eisenhart-Rothe, M. Hudelmaier, D. Wilson, and F. Eckstein. Use of novel interactive input devices for segmentation of articular carti-

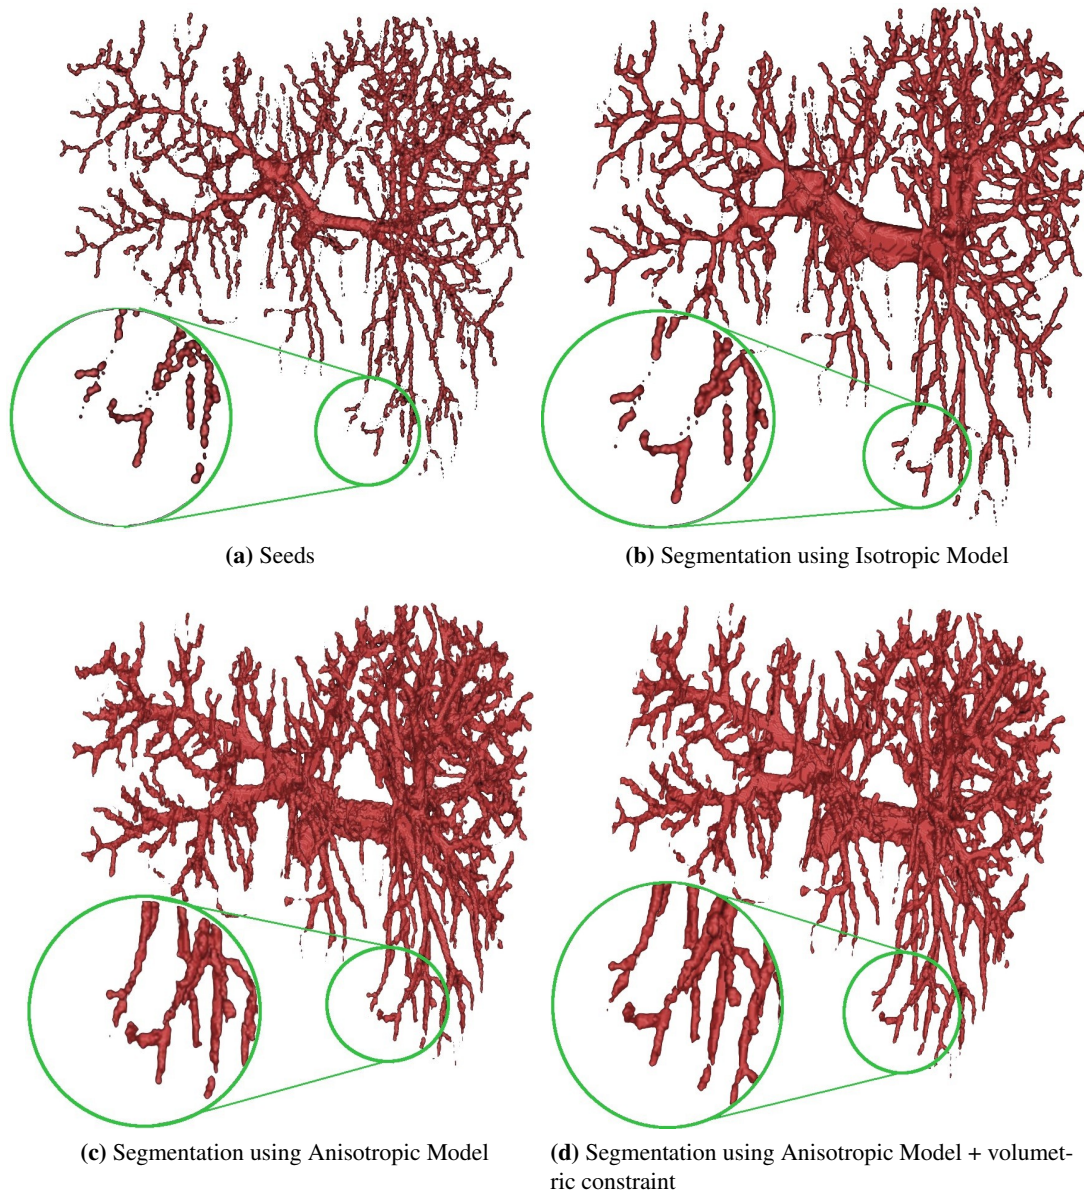


Figure 7: Segmentation of liver vessels: The impact of using directional edge information and a global volume constraint is shown in the magnification

- lage from magnetic resonance images. *Osteoarthritis and Cartilage*, 13(1):48 – 53, 2005. 5
- [19] S. Olabbariaga and A. Smeulders. Interaction in the segmentation of medical images: A survey. *Medical Image Analysis*, 5:127–142, 2001. 1
- [20] C. Olsson, M. Byröd, N. Overgaard, and F. Kahl. Extending continuous cuts: Anisotropic metrics and expansion moves. In *International Conference on Computer Vision*, 2009. 2
- [21] T. Pock, D. Cremers, H. Bischof, and A. Chambolle. An algorithm for minimizing the mumford-shah functional. In *IEEE International Conference on Computer Vision (ICCV)*, 2009. 4
- [22] M. Unger, T. Pock, W. Trobin, D. Cremers, and H. Bischof. TVSeg - interactive total variation based image segmentation. In *Proceedings of British Machine Vision Conference*, 2008. 2
- [23] P. Yushkevich, J. Piven, C. Hazlett, G. Smith, H. Smith, R. Ho, S. Ho, J. Gee, and G. Gerig. User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. *Neuroimage*, 31:1116–1128, 2006. 4