Eurographics Italian Chapter Conference (2008)
V. Scarano, R. De Chiara, and U. Erra (Editors)

Medical Image Segmentation using Level Sets

Gaetano Impoco

Dipartimento di Matematica e Informatica, Università di Catania, Italy
e-mail: impoco@dmi.unict.it

Abstract

Computer-aided diagnosis for pre-operative planning and post-operative outcome evaluation is widely considered an important topic for next-generation surgery. 3D models of the patients' anatomical structures can be highly valuable in this context. The accuracy of these models is strongly dependent on the classification and segmentation algorithms acting at the very first stage of the modelling chain. A promising class of segmentation algorithms is related to level set *methods. Here, we briefly review some applications of level sets to medical image segmentation.*

Categories and Subject Descriptors (according to ACM CCS): I.4.6 [Image Processing and Computer Vision]: Segmentation, Level Sets, 3D, MRI, CT, Medical Imaging

1. Introduction

The rapid technological advances in digital medical imaging devices have brought the attention of various medical communities to computer-assisted diagnosis.

One of the trends of modern surgery is computer-aided pre-operative planning and post-operative outcome evaluation. Digital models of the involved organs or tissues can be of great value both to predict and to evaluate the outcome of a surgical intervention. In particular, planning and evaluation often involve the creation of a model from patient-specific data, acquired by means of CT, MRI, and so on. 3D models of human body parts can be also useful for selecting perfectly-fitting prostheses e.g., for knee replacement [BBFI06] (using FEM stress analysis) and plastic surgery [FIG*06] (by means of ad-hoc geometric measurements).

Acquiring volumetric data of internal organs by means of CT, MRI, or PET, is common clinical practice both for diagnosis and pre-operative planning. Although this information could be fruitfully exploited for quantitative measurements, currently it is mainly employed qualitatively by surgeons. This is mostly due to the lack of robust and reliable systems to build precise 3D models of the structures of interest. A large variety of segmentation methods have been developed for medical image processing and segmentation. Nonetheless, ad-hoc solutions are often preferred to properly detect complex structures, such as vessels, organs, or skeletal structures.

The main difficulties with CT, MRI, or PET images are related to acquisition noise, anatomical peculiarities, and low resolution. Noise reduction is probably the easiest to cope with, since lots of noise reduction methods have been developed for 2D images. However, care must be taken to extend these algorithms to 3D anisotropic grids. Resolution causes the *partial volume (PV) effect* i.e., a single voxel may contain a number of different tissues. A common problem derived from PV effect is the difficulty to separate bones in an articulation. Resolution is going to increase in the near future due to the fast technical advances in this field. Anyway, the amount of radiation each patient can safely absorb will pose a limit for most real cases. Finally, anatomical features may vary in different patients and even at different parts of the body of a patient (e.g., the density of a femur is remarkably lower close to the articulation than in its central area).

All these problems rule out many traditional segmentation algorithms. On the other hand, anatomy can be a valuable information to guide segmentation. Hence, geometric features and shape priors can be fruitfully exploited for segmentation.

Active Contours, also called *snakes* [KWT88], can easily handle shape-driven segmentation by explicitly introducing geometric constraints, modelled as curvature terms in a curve evolution setting. A binary segmentation is obtained by separating the inner and the outer side of the evolving

© The Eurographics Association 2008.

delivered by EUROGRAPHICS DIGITAL LIBRARY www.eg.org diglib.eg.org

contour. The evolution of the contour is guided by differential equation, encompassing different geometric and image contributions. The most common terms of the evolution equation are a curvature term to smooth out cusps, an advection term (i.e., an inflation/deflation balloon force), and an image gradient term to force the contour to stick to image edges.

Snakes have been very popular and extensively studied in the last decade, especially for medical image segmentation. Despite their success, they have two serious drawbacks. First, topological variations occurring during the evolution, such as region merging and splitting, are not handled easily. When designing new evolution equations one must account of a number of tricky special cases. Second, the contour can overlap and fold during evolution (it is no more a simple curve), resulting in unwanted unnatural effects. Again, avoiding folding can be hard and result in a time consuming implementation.

In order to circumvent these problems, the *Level Set* paradigm was introduced by Osher and Sethian [OS88] and rapidly showed its potential for many applications. Among them, medical image segmentation. In the following sections, we will give a short introduction to the level set formulation and survey some of the most interesting medical image segmentation algorithms since their introduction.

A joint project involving our lab and numerical analysts of our Department has recently started to investigate the topics addressed in this paper.

2. Level Set Basics

Most of the problems with active contours are due to parametrisation. Basically, in order to evolve a moving contour one should follow the path of each (infinitesimal) particle of the curve (Langrangian approach). Discretising the evolving contour is far from being straightforward. One could fix some knots at fixed equal distances along the curve, and evolve the position of these knots. However, as pointed out in [Set97], there is no means to detect if their relative position along the curve will switch at some time during the evolution, without re-parameterising the contour. This inversion would then cause the curve folding over itself. This problem could be solved by stopping the evolution and discretising again the curve. However, doing this can be computationally burdensome, especially for surfaces in three dimensional space.

An even more problematic case is topology change due to curve merging and splitting. In order to illustrate this problem, let us take a simple example of two growing circular contours in the same plane [Set97]. At some time during evolution, they will touch at one point and then merge. However, since time steps are discretised we may not be able to track the exact time when the contours meet. Thus, it will be hard to find and remove the exceeding knots lying inside the merged contour, and restore the correct knot ordering. This problem is even more complex in three dimensions.

Rather than following the evolution of contour particles, the level set formulation tracks the time when the contour crosses each point in space (Eulerian approach) [OF02]. An extra dimension (time) is added to the problem, so that the evolving contour is a section of a higher dimensional embedding (a *level set*) at a given time. Doing this way, topology changes of the contour are handled easily since no topology change occurs in the embedding. Thus, the choice of grid step is not crucial. Moreover, the problem of knots order switching is avoided since the space is sampled, rather than the evolving contour itself. Hence, none of the problems cited above will come up and a discretisation as simple as a regular grid will suffice. Other important advantages of level sets with respect to active contours are that the mathematical approach is independent from problem dimensionality and that anisotropy of the grid can be easily dealt with. This is particularly useful for medical image segmentation.

Before describing how level sets can be used for segmentation, we briefly introduce some fundamental mathematical notions. Our discussion is based on planar curves; however, the same formulation can be used for higher dimensional problems. Let $C \in \mathbb{R}^2$ be a closed curve. Active contours parameterise the curve as $C(p) \equiv C(x(p), y(p))$, where $p \in [0,1]$ is a parameter and $C(0) = C(1)$. Level sets instead use an implicit representation. Namely, an embedding function $\Phi : \mathbb{R}^3 \to \mathbb{R}$ such that $C_t = \{(x,y) | \Phi(x,y,t) = k\}$, with $k \in \mathbb{R}$ arbitrary (usually set to zero). Although there are several choices for Φ, most applications use signed distance functions since the resulting math is simplified and numerical approximations are more stable [OF02]. Setting to zero the time derivative of $\Phi(x,y,t) = k$ we get

$$\frac{\partial \Phi}{\partial t} = -\nabla\Phi \cdot \frac{d\vec{x}}{dt} = -\nabla\Phi \cdot \mathbf{F} \tag{1}$$

where $\mathbf{F}$ is a function encompassing the partial derivatives of Φ evaluated at $\vec{x} \equiv (x,y)$. $\mathbf{F}$ acts as an external force, driving the evolution of the contour. Equation 1 is referred to as the *level set equation*. Note that if $\mathbf{F}$ is constant, the term $\nabla\Phi \cdot \mathbf{F}$ (*advection* or *convection* term) acts as a contraction/expansion balloon force. If we focus on motion along the normals to the contour (i.e., the tangential component is zero), Equation 1 can be rewritten as

$$\frac{\partial \Phi}{\partial t} = -V_n |\nabla\Phi| \tag{2}$$

where V_n is the *normal velocity* i.e., the component of the velocity in the normal direction. This is the case of motion by mean curvature: $V_n = c \cdot \kappa$, where κ is the mean curvature and c is a constant. Notice that $c \cdot \kappa |\nabla\Phi|$ represents an internal force, since it depends only on the contour. Evolving the contour using a curvature-driven flow, the high-curvature features of the contour move significantly faster. It can be shown [Gra87] that, under motion by mean curvature, any simple closed curve evolves towards a circular shape and

© The Eurographics Association 2008.

than collapses to a point. Hence, $c \cdot \kappa |\nabla \Phi|$ acts as a regularisation term. Putting together the balloon force and the curvature-driven flow, the level set equation becomes

$$\frac{\partial \Phi}{\partial t} = (1 - c \cdot \kappa)|\nabla \Phi| \tag{3}$$

Since here we are not interested in implementation details, in the discussion above we did not mention any approximation scheme for the numerical solution of the level set equation. The interested reader is referred to [OF02].

3. Level Sets Applied to Medical Image Segmentation

Level sets can be usefully employed for image segmentation by adding an image-dependent external force to Equation 3. Since regions are bounded by edges, diffusion should be arrested in correspondence to image gradients. Given an image I, an edge detector can be defined as a positive decreasing function, $\Psi_I(\vec{z})$, of the image gradient ∇I, such that $\Psi_I(\vec{z}) \to 0$ for $\vec{z} \to \infty$. A common image-dependent term can thus be defined as

$$\Psi_I(\nabla I(\vec{x})) = \frac{1}{1 + |\nabla G_\sigma * \nabla I(\vec{x})|^p} \tag{4}$$

where G_σ is a Gaussian of variance σ, $*$ is the convolution operator, and $p \geq 1$ is usually 1 or 2. The effect of this term is to force the evolution to slow down and stop when the contour is close to intensity gradients. Plugging Equation 4 into the embedding function and setting to zero its derivative as in Equation 1, we get

$$\frac{\partial \Phi}{\partial t} = \Psi_I(\nabla I)(1 - c \cdot \kappa)|\nabla \Phi| + \Psi_I(\nabla I) \cdot \nabla \Phi \tag{5}$$

Equation 5 is the level set formulation (see [YKK*97] and [ZOF01]) of the *geodesic active contours*, derived by Caselles et al. [CKS97] and Kichenassamy et al. [KKO*96]. It is a standard formulation used in most segmentation algorithms.

An alternative formulation was proposed in [CV99] and [CV01] for binary segmentation, and later extended to cope with multi-region segmentations. Let $\mu_1(\Phi)$ and $\mu_2(\Phi)$ be the mean intensities of the inner and outer regions of a contour, C, in the embedding Φ. Then, the level set function can be written as

$$\frac{\partial \Phi}{\partial t} = \delta_\varepsilon(\Phi)\left[\mu \nabla \cdot \left(\frac{\nabla \Phi}{|\nabla \Phi|}\right) - \nu - \gamma_1 (I - \mu_1)^2 - \gamma_2 (I - \mu_2)^2\right] \tag{6}$$

where μ is the strength of the smoothness (curvature) term, γ_1 and γ_2 weight the internal coherence of the inner and outer regions, and ν is a balloon force. The function $\delta_\varepsilon(\Phi)$ is a smooth approximation to the delta function, where ε controls its smoothness and should be related to grid step size. A common choice for these parameters is $\mu > 0$, $\gamma_1 = \gamma_2 = 1$, and $\nu = 0$. Two important observations must be made here. First, even if only two regions can be constructed they can be disconnected into a number of fine-scale components. Second, the mean values μ_1 and μ_2 are global image operators. This is a remarkable difference between Equation 5 and Equation 6: While the former acts locally, the latter must compute a global operator at image level. In the following, global methods will be presented that compute global operators on the embedding that, as such, require Φ to be defined in the whole domain.

3.1. Segmentation of Blood Vessels and Codimension Two Objects

Blood vessels in a CT or MRI scan can be represented as tubular structures in a 3D volume. Great care must be taken using the level set equation to segment these structures for two reasons. First, the smoothness term in Equation 5 is usually taken as the mean or as the larger principal curvature of the surface. In contrast, when segmenting tubular structures we require smoothness along the tube, orthogonally to its section. The principal curvature generally captures the curvature of the tube section which is greater. A simple solution is to take the smaller of the two principal curvatures as the curvature term in the level set equation.

A more subtle problem is related to tubes with infinitesimal cross-section i.e., curves in space. We might be interested in these structures to describe, for example, the skeleton of a blood vessel. In this case, the points on the curve are singular since the distance function is null on the curve and positive elsewhere. This problem is avoided for surfaces due to the sign change of the distance function from the inside to the outside of the surface. In the case of curves in space, no inside or outside is defined; the same is true for any object with codimension greater than one (i.e., an object defined by k variables embedded in a $\mathbb{R}^d$ space, with $k < d - 1$).

The solution proposed in [LFG*01] is to embed the curve, C, into distance field ϕ such that each isolevel is a thin tube around C. The curve C itself is the zero level set of this embedding. ϕ is then evolved using a special evolution equation, which we do not report here. After convergence, the final curve is obtained extracting the zero level set of the embedding.

A different approach is presented by van Bemmel et al. in [vBSVN03] for segmentation of blood vessels in angiography. They are not interested in curves in space, thus Equation 5 for codimension one objects is used to evolve the contour. However, they drop the last term and, most important, they use a strongly different definition for the image term. Ψ_I is computed as the product of three terms: a gradient term Ψ_{grad}, an intensity term Ψ_{int}, and a structure term Ψ_{vessel}. Ψ_{grad} pushes the contour towards gradient edges, in the same spirit as in Equation 4, using an exponential function of the image gradient. Ψ_{int} is the normalised difference of two Normal distributions fitting vessels and background intensities in the data, respectively. This term captures the bimodal behaviour of angiograms. The last term, Ψ_{vessel}, is

© The Eurographics Association 2008.

a geometric term, encoding prior knowledge about the tubular geometry of the structures of interest. Three components are computed to discriminate structures, respectively, with respect to round/flat sections, tubular/bubble shapes, and intensity variability (high variance is expected inside vessels). These components are computed from the eigenvalues, λ_1, λ_2, and λ_3, of the image Hessian. The two largest eigenvalues, λ_1, λ_2, are related to the highest variations, thus represent the diameters of vessel cross-section. λ_3 is related to the orthogonal direction. Hence, for example, if $|\lambda_1| \approx |\lambda_2|$ the tube section is round. Similarly, $|\lambda_3| \ll \min(|\lambda_1|, |\lambda_2|)$ implies a tubular structure.

The same formulation of the level set equation is used in [MN04]. Equation 5 is used without the last term. Again, the curvature term is set as the minor of the two principal curvatures. In this case, however, the image term Ψ_I is defined as the normalised difference of the number of voxels being misclassified, respectively, as background or as vessel. Two Gaussian distribution, fitted to the data by means of the EM algorithm, are used to estimate the total classification error for background and vessel regions. This is similar to the intensity term, Ψ_{int}, in the previous approach. A small value indicates a balance between misclassification errors. Hence, evolution slows down and stops when this equilibrium is reached. The problem of extracting tubes with infinitesimal cross-sections (skeleton of the vessels) is solved by labelling inner (vessel) and outer (background) voxels and then using morphological thinning on vessel voxels. Rather than running the evolution on the whole volume, *volumes of interest* (VOIs) are selected and the process repeated for each VOI separately. A user-selected seed point is used to initialise the first VOI. After the evolution converges, the voxels are classified as vessel or background. The skeleton of the vessel segment is then extracted, one endpoint corresponding to the seed point. The other endpoint is used as a seed to initialise a new VOI. After a new segment of the vessel has been classified, it is combined using a simple OR with the parts classified in previous steps. The work was later extended in [MVvL*06], by adding a simple morphological bone masking operation which employs two registered scans with and without contrast fluid, respectively.

As a last example, we summarise the work presented in [NYT04]. Equation 5 is used without the last term. The image term, Ψ_I, models the closeness of a voxel to a vessel border. During the evolution, a value ε is computed for each voxel, representing the number of vessel voxels lying in a ball of radius r centred in the voxel. ε is high for inner voxels and low for border cells. For each voxel, Ψ_I is set as the sum of the values of ε for all neighbouring voxels in a ball of radius r. To see why this measure can distinguish vessel voxels from potential leaks when r is close to the expected vessel radius, let us consider two border voxels, one lying on a vessel border, the other on the border of a leak. All voxels inside a vessel of radius r or less have at least one neighbouring border voxel. Thus, ε is always lower than the size of the neighbourhood. Conversely, for a structure wider than r the entire neighbourhood of a voxel can lie inside the structure. Hence, ε equals the neighbourhood size. In the first case, Ψ_I is low since ε is low for all neighbours of the current voxel. In the second case, Ψ_I takes a higher value since most of the neighbours of the current voxel have high ε values. From this discussion, it should be clear that Ψ_I is higher for voxels lying in structures wider than r voxels. Hence, if we expect that vessels have a maximum radius r, vessel voxels can be discriminated checking the value of Ψ_I.

3.2. Introducing Priors into Segmentation

When the shape of the objects to be segmented is known beforehand, shape priors (i.e., a-priori shapes used when some salient features of the final curve are known in advance) can be employed to improve segmentation. One of such cases is medical images: the shape of anatomical structures is well known. There are several ways of introducing prior knowledge into the evolution equation, such as local or global statistical analysis, shape priors, and implicit shape functions. One such example is the last method introduced in Section 3.1. Here, prior knowledge about vessel shape and width is embedded into the image term as an implicit function. Geometric priors are used in the algorithm by van Bemmel et al., described in the same section.

Methods employing shape priors are easy to use due to the naturalness of their interface. Basically, drawing a shape template will mostly suffice. However, defining templates invariant under rotation, translation, and scaling may be challenging. Invariance is addressed explicitly in [CTT*02]. A planar curve C is defined as a rotated, scaled, and translated version of a template curve, C^*. The evolution equation contains a shape term derived from a distance function that maps each point of C onto the closest point of C^*. At each step of the level set evolution, the rotation, translation, and scaling parameters are computed by minimising this distance function, in order to adapt the template to the current segmentation. Level set evolution is used for minimisation. After the parameters have been optimised, they are plugged into the curve evolution equation to update the evolving curve, C. Notice that this algorithm requires a global minimisation of the parameters at each step of the evolution. It can thus be prohibitively time consuming. In order to learn C^* from a training set, a similarity measure is defined between manually-drawn curves, based on the percentage of overlapping internal area. All curves in the training set are aligned by minimising this measure. The template, C^*, is set as the mean of the aligned shapes.

The work in [CSS03] is motivated by the need to extract expected objects of interest without excluding other objects i.e., the template should be used only when needed. The idea is to automatically estimate a function, L, encoding the similarity between portions of the evolving contour and the prior shape. The evolution equation is designed such that the prior

© The Eurographics Association 2008.

shape influences only the portions of the contour which show a high degree of similarity to the template. The evolution Equation 6 is used, enriched with a shape term given by the squared difference between the current embedding, Φ, and the template, Φ_0, scaled by the value of the template similarity function, L. Then, the evolution equations of Φ and of L are simultaneously minimised. L is evolved such that its influence on the evolution of Φ is strong when the shape is similar to the template, and negligible when this is not the case. Notice that this method is not invariant with respect to any motion or scaling parameter. This is due to the trivial definition of the shape term as a simple difference between Φ and Φ_0.

Similarly to the previous method, Equation 6 together with a shape term is used in [CZ05]. The shape term is similar to the corresponding term of the previous method. Further, another term is used to force the prior shape to give a good segmentation of the image. That is, the influence of the template is low if it does not fit the image. The translation, rotation, and scaling parameters of the prior shape are simultaneously estimated similarly to the first method in this section. Thus, this method is computationally intensive. The effect is to fit the prior shape to the image, if possible, and then evolve the contour towards the template.

Statistical priors are used in [LFGW00]. A Markov network is used to relate the value of each cell to its neighbourhood. The posterior probability has an image term and a regularisation term. The image term is modelled as a weighted mean of the distance functions from the contours in the training set. The regularisation term relates a cell with its neighbouring cells. It is composed by a tangent and a normal contribution, under the hypothesis of statistical independence, related to curvature and to linearity, respectively. The normal term is modelled using central differences, while the tangent term is computed as the deviation of the curvature from the central difference in the tangent direction. The evolution equation minimises the log posterior probability, obtained as the composition of the three terms above.

Another method [LGF00] extends Equation 5 by adding a statistical shape term, computed as the difference between the evolving curve and the template, suitably rotated and translated (no scaling is allowed). Shape and pose parameters are statistically estimated using a maximum a posteriori approach. Three terms are considered. The first one is an inside term which decreases as more and more cells fall outside the estimated final template curve. The second term computes the deviation of the image profile along the curve gradient from a Gaussian fitted to the data along the template curve. The third term relates the shape parameters to possible prior shapes. Finally, a uniform distribution over pose parameters is assumed.

4. Conclusions

We have reviewed some of the most interesting level set approaches to medical image segmentation. We chose an informal approach to focus on ideas rather than on the formalisms used. We think that level sets with statistical and shape priors have the potential to overcome many of the limitations of classical segmentation algorithms. We will thus devote our research efforts to these methods.

5. Acknowledgements

Many thanks go to Prof. Giovanni Gallo for his precious advice.

References

[BBFI06] BATTIATO S., BOSCO C., FARINELLA G. M., IMPOCO G.: 3D CT segmentation for clinical evaluation of knee prosthesis operations. In *Proceedings of the Fourth Eurographics Italian Chapter Conference* (22-24 February 2006).

[CKS97] CASELLES V., KIMMEL R., SAPIRO G.: Geodesic active contours. *International Journal on Computer Vision 22* (1997), 61–79.

[CSS03] CREMERS D., SOCHEN N. A., SCHNÖRR C.: Towards recognition-based variational segmentation using shape priors and dynamic labeling. In *Scale Space Methods in Computer Vision* (2003), vol. 2695/2003 of *Lecture Notes in Computer Science*, pp. 388–400.

[CTT*02] CHEN Y., TAGARE H. D., THIRUVENKADAM S., HUANG F., WILSON D., GOPINATH K. S., BRIGGS R. W., GEISER E. A.: Using prior shapes in geometric active contours in a variational framework. *International Journal on Computer Vision 50*, 3 (2002), 315–328.

[CV99] CHAN T., VESE L.: An active contour model without edges. *Lecture Notes in Computer Science (LNCS) 1682/1999* (1999), 141–151.

[CV01] CHAN T., VESE L.: Active contours without edges. *IEEE Transactions on Image Processing 10* (2001), 266–277.

[CZ05] CHAN T., ZHU W.: Level set based shape prior segmentation. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05) - Volume 2* (Washington, DC, USA, 2005), IEEE Computer Society, pp. 1164–1170.

[FIG*06] FARINELLA G. M., IMPOCO G., GALLO G., SPOTO S., CATANUTO G., NAVA M.: Objective outcome evaluation of breast surgery. *Lecture Notes in Computer Science 4190/2006* (2006).

[Gra87] GRAYSON M.: The heat equation shrinks embedded plane curves to round points. *Journal of Differential Geometry 26*, 285 (1987).

© The Eurographics Association 2008.

[KKO*96] KICHENASSAMY S., KUMAR A., OLVER P., TANNENBAUM A., YEZZI A.: Conformal curvature flows: from phase transitions to active vision. *Archive for Rational Mechanics and Analysis 134* (1996), 275–301.

[KWT88] KASS M., WITKIN A., TERZOPOULOS D.: Snakes: Active contour models. *International Journal of Computer Vision 1*, 4 (January 1988), 321–331.

[LFG*01] LORIGO L. M., FAUGERAS O. D., GRIMSON W. E. L., KERIVEN R., KIKINIS R., NABAVI A., WESTIN C.-F.: CURVES: Curve evolution for vessel segmentation. *Medical Image Analysis 5*, 3 (September 2001), 195–206.

[LFGW00] LEVENTON M. E., FAUGERAS O., GRIMSON W. E. L., WELLS W. M.: Level set based segmentation with intensity and curvature priors. In *IEEE Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA '00)* (Washington, DC, USA, June 2000), IEEE Computer Society, pp. 4–11.

[LGF00] LEVENTON M. E., GRIMSON W. E. L., FAUGERAS O.: Statistical shape influence in geodesic active contours. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR00)* (2000), pp. 316–323.

[MN04] MANNIESING R., NIESSEN W.: Local speed functions in level set based vessel segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)* (Berlin, September 2004), Barillot C., Haynor D. R., Hellier P., (Eds.), vol. 3216 of *Lecture Notes in Computer Science*, pp. 475–482.

[MVvL*06] MANNIESING R., VELTHUIS B. K., VAN LEEUWEN M. S., VAN DER SCHAAF I. C., VAN LAAR P. J., NIESSEN W. J.: Level set based cerebral vasculature segmentation and diameter quantification in CT angiography. *Medical Image Analysis 10*, 2 (2006), 200–214.

[NYT04] NAIN D., YEZZI A. J., TURK G.: Vessel segmentation using a shape driven flow. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)* (2004), pp. 51–59.

[OF02] OSHER S., FEDKIW R. P.: *Level Set Methods and Dynamic Implicit Surfaces*. Springer, October 2002.

[OS88] OSHER S., SETHIAN J. A.: Fronts propagating with curvature-dependent speed: algorithms based on Hamilton-Jacobi formulations. *Journal of Computational Physics 79*, 1 (1988), 12–49.

[Set97] SETHIAN J. A.: Level set methods: An act of violence. *American Scientist 85*, 3 (May-June 1997).

[vBSVN03] VAN BEMMEL C., SPREEUWERS L., VIERGEVER M., NIESSEN W.: Level-set based artery-vein separation in blood pool agent CE-MR angiograms. *IEEE Transactions on Medical Imaging 22*, 10 (October 2003), 1224–1234.

[YKK*97] YEZZI A., KICHENASSAMY S., KUMAR A., OLVER P., TANNENBAUM A.: A geometric snake model for segmentation of medical imagery. *IEEE Transactions On Medical Imaging 16*, 2 (April 1997).

[ZOF01] ZHAO H.-K., OSHER S., FEDKIW R.: Fast surface reconstruction using the level set method. In 1st *IEEE Workshop on Variational and Level Set Methods,* 8th *ICCV* (2001), pp. 194–2002.

© The Eurographics Association 2008.