



HAL
open science

Heart segmentation with an iterative Chan-Vese algorithm

Olivier Rousseau, Yves Bourgault

► **To cite this version:**

Olivier Rousseau, Yves Bourgault. Heart segmentation with an iterative Chan-Vese algorithm. 2008.
hal-00403627v2

HAL Id: hal-00403627

<https://hal.science/hal-00403627v2>

Preprint submitted on 16 Aug 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

HEART SEGMENTATION WITH AN ITERATIVE CHAN-VESE ALGORITHM

OLIVIER ROUSSEAU, YVES BOURGAULT

ABSTRACT. This paper presents 2D and 3D applications of the Chan-Vese model to heart and trachea segmentation. We improved the multi-phase Chan-Vese model by introducing an iterative method, by choosing an appropriate L^1 fidelity term as well as an efficient and prior free initial condition. For 3D applications, the algorithm is parallelized in order to speed up the computations. We provide extensive information on computational details, on the convergence times and on the quality of segmentations obtained. The results of the segmentations are then meshed to be used for finite element simulations.

INTRODUCTION

The recent increase of computers capacities now allows for realistic simulations of human organ physiology. Mathematical modeling of organ functions opens a wide range of research, allowing diagnostics and understanding of malfunctions and diseases. Such simulations usually require a mesh of the given organ. However, most computations are made on meshes of idealized geometries and there is a real lack of accurate 3D models. Realistic geometries should be extracted from medical images, this is known as the segmentation process.

The goal of this work is to segment the heart muscle from high resolution CT scans of the thorax and to produce meshes that are adequate for numerical simulations in electro-physiology.

Most existing methods for heart segmentation involve a prior knowledge of the heart's shape (see for example [29, 12]). We intend to use a modified version of the Chan-Vese model [9], also known as *Active Contours without Edges* to segment the heart, since it has no geometrical or topological *a priori*. This method has been successfully used for brain segmentation [28, 10]. Brain images usually have nice contrasts between gray and white matter. To our knowledge, this method has not been attempted yet for segmenting 3D scans of the heart and lungs. This is a real challenge since the images are more diversified. They contain many objects of similar grey levels that the method should separate.

The main contributions of this paper are:

- (1) Application of the Chan-Vese algorithm to trachea and heart segmentation.
- (2) Introduction of an Iterative version of the Chan-Vese algorithm to replace multi-phase segmentation.
- (3) Analysis of the convergence and efficiency of the Chan-Vese algorithm under different initial conditions and different fidelity terms.

O. Rousseau and Y. Bourgault are with the Department of Mathematics and Statistics, University of Ottawa, Ontario, Canada. They can be joined at orous097@uottawa.ca and ybourg@uottawa.ca.

- (4) Analysis of parallelization of the Chan-Vese algorithm to fit needs of large applications.
- (5) Creation of meshes from the segmentation results.

In the first section, the Mumford-Shah functional and the Chan-Vese model are presented as well as the numerical methods for solving the latter. In section 2, the improvements made to the original Chan-Vese model are explained and justified. Sections 3 and 4 respectively present 2D and 3D applications of this modified Chan-Vese algorithm.

1. BACKGROUND

1.1. The Mumford-Shah energy and the Chan-Vese model. An image can be interpreted as a function

$$g : \Omega \longrightarrow \mathbb{R},$$

where Ω is some region of \mathbb{R}^n , typically a square or a cube. In the segmentation process, the goal is to split the image g into its constituting objects. Mumford and Shah proposed to minimize the following functional

$$(1) \quad E_{MS}(u, K) = \int_{\Omega} |\nabla u|^2 dx + \lambda \int_{\Omega \setminus K} |g - u|^2 dx + \mu \mathcal{H}^{N-1}(K)$$

over pairs (u, K) . K is a compact subset of Ω representing edges of objects in g , and $u \in H^1(\Omega \setminus K)$ is the intensity of the image. This intensity varies smoothly inside the connected components of $\Omega \setminus K$ [20]. The middle term in the equation is called the fidelity term, while the combination of the two others form the regularity terms that do not depend on the underlying image g . The weights μ and λ should be adjusted in accordance with the noise level of the image to be segmented.

It is a well known result that minimizers of the Mumford-Shah energy exist (see for example [4]). It is a hard task to directly find the actual minima of the functional. To achieve this, one can approximate the Mumford-Shah functional by a sequence of elliptic functionals easier to solve [5, 6, 7]. An other approach is to minimize E_{MS} over a restricted domain. One example is the Chan-Vese model. It seeks for a minimum over functions that take only two values c_1 and c_2 [9]. Such functions can be written as $u = c_1 \mathbb{1}_F + c_2(1 - \mathbb{1}_F)$ for some $F \subseteq \Omega$. The Mumford-Shah energy then rewrites as

$$(2) \quad E_{CV}(F, c_1, c_2) = \int_F |g - c_1|^2 dx + \int_{\Omega \setminus F} |g - c_2|^2 dx + \mathcal{H}^{N-1}(\partial F \cap \Omega).$$

The set F can be described via a level set function $\phi : \Omega \rightarrow \mathbb{R}$, that is $F = \{\phi \leq 0\}$. Then the Level Set formulation [24, 9] of the problem 2 is to find ϕ that minimizes

$$(3) \quad E_{CV}(\phi, c_1, c_2) = \lambda \int_{\Omega} |g - c_1|^2 (1 - H(\phi)) dx + \lambda \int_{\Omega} |g - c_2|^2 H(\phi) dx + \mu \int_{\Omega} |DH(\phi)| dx,$$

where $H(\cdot)$ stands for the Heaviside function. $|DH(\phi)| = \delta(\phi)|\nabla\phi|$ is the derivative of $H(\phi)$ in the sense of distributions. This is a Dirac measure with support on the discontinuity set $\{x : \phi(x) = 0\}$ of u . The two-phase image then rewrites as $u = c_1(1 - H(\phi)) + c_2H(\phi)$.

We can deduce that if (ϕ, c_1, c_2) is a minimum of E_{CV} , then it is one of the many solutions of the Euler-Lagrange equation. To avoid difficulties related to the

non-uniqueness of ϕ , a time dependence is added and the following initial problem is solved:

$$(4) \quad \begin{cases} \frac{\partial \phi}{\partial t} = \delta(\phi) \left[\mu \operatorname{div} \left(\frac{\nabla \phi}{|\nabla \phi|} \right) + \lambda ((g - c_1)^2 - (g - c_2)^2) \right] \\ \phi(x, 0) = \phi_0(x), \\ \frac{\partial \phi}{\partial n} = 0 \quad \text{on } \partial\Omega, \end{cases}$$

for some initial condition ϕ_0 which describes an initial curve $C_0 = \{\phi_0 = 0\}$. The values c_1 and c_2 are also given by

$$(5) \quad c_1 = \frac{\int_{\Omega} g(1 - H(\phi)) dx}{\int_{\Omega} (1 - H(\phi)) dx}, \quad c_2 = \frac{\int_{\Omega} gH(\phi) dx}{\int_{\Omega} H(\phi) dx},$$

That is, c_1 and c_2 are mean values of g inside the regions $\{\phi < 0\}$ and $\{\phi \geq 0\}$ respectively.

While solved with a time-stepping scheme, the problem 4 can be interpreted as a gradient method applied to the minimization of the functional in 3. Any steady state of 4 will satisfy the Euler-Lagrange equation.

In order to get a splitting of the image into more than 2 phases (multi-phase segmentation), Chan and Vese proposed to have several level set functions ϕ_i [28]. A set of n level set functions can describe up to 2^n different regions of Ω . In this case n coupled PDE's must be solved simultaneously. An other approach is to use several level curves of the level set function [10]. In this case the number of regions has to be chosen in advance. This method does not allow to have junctions of three different regions as level curves are parallel to each other.

1.2. Numerics. The equation 4 can be solved via standard finite difference scheme ([9], [23]) using a C^∞ regularization δ_ϵ of δ . For example, in 2D the first term of $\operatorname{div} \left(\frac{\nabla \phi}{|\nabla \phi|} \right)$ can be discretized as

$$(6) \quad \left(\frac{\partial}{\partial x} \frac{\phi_x}{\sqrt{\phi_x^2 + \phi_y^2}} \right)_{i,j} \approx D_x^- \left[\frac{D_x^+ \phi_{i,j}}{\sqrt{(D_x^+ \phi_{i,j})^2 + \left(\frac{D_y^0 \phi_{i,j} + D_y^0 \phi_{i+1,j}}{2} \right)^2 + \epsilon^2}} \right].$$

$D_x^+ \phi$, D_x^0 , D_x^- respectively stand for forward, centered and backward finite difference approximations. Hence an explicit discretization of equation 4 can be written as

$$(7) \quad \begin{aligned} \phi_{i,j}^{n+1} = & \phi_{i,j}^n + \Delta t \delta_\epsilon(\phi_{i,j}) [d_1(\phi_{i+1,j}^n - \phi_{i,j}^n) + d_2(\phi_{i,j}^n - \phi_{i-1,j}^n) \\ & + d_3(\phi_{i,j+1}^n - \phi_{i,j}^n) + d_4(\phi_{i,j}^n - \phi_{i,j-1}^n) \\ & + \lambda ((g - c_1)^2 - (g - c_2)^2)]. \end{aligned}$$

This is the form of a non linear diffusion with coefficients

$$\begin{aligned}
 d_1 &= \frac{1}{h_x^2} \frac{\mu}{\sqrt{(D_x^+ \phi_{i,j})^2 + \left(\frac{D_y^0 \phi_{i,j} + D_y^0 \phi_{i+1,j}}{2}\right)^2 + \varepsilon^2}} \\
 d_2 &= \frac{1}{h_x^2} \frac{\mu}{\sqrt{(D_x^- \phi_{i,j})^2 + \left(\frac{D_y^0 \phi_{i,j} + D_y^0 \phi_{i-1,j}}{2}\right)^2 + \varepsilon^2}} \\
 d_3 &= \frac{1}{h_y^2} \frac{\mu}{\sqrt{\left(\frac{D_x^0 \phi_{i,j} + D_x^0 \phi_{i,j+1}}{2}\right)^2 + (D_y^+ \phi_{i,j})^2 + \varepsilon^2}} \\
 d_4 &= \frac{1}{h_y^2} \frac{\mu}{\sqrt{\left(\frac{D_x^0 \phi_{i,j} + D_x^0 \phi_{i,j-1}}{2}\right)^2 + (D_y^- \phi_{i,j})^2 + \varepsilon^2}}
 \end{aligned}
 \tag{8}$$

However, for an explicit time discretization, the time step constraint for stability may be quite restrictive, leading to a large number of time steps to reach a steady solution.

It has been noticed that replacing $\delta(\phi)$ by $|\nabla\phi|$ increases the stability of the PDE [19], yielding

$$\phi_t = |\nabla\phi| \left[\mu \operatorname{div} \left(\frac{\nabla\phi}{|\nabla\phi|} \right) + \lambda ((g - c_1)^2 - (g - c_2)^2) \right],
 \tag{9}$$

which can be interpreted as a motion of the interface in the normal direction with speed $V_n = -\mu \operatorname{div} \left(\frac{\nabla\phi}{|\nabla\phi|} \right) - \lambda ((g - c_1)^2 + (g - c_2)^2)$. Note that equation 9 does not admit any steady state, although the interface $\{\phi = 0\}$ converges to the zero level set of the steady state of equation 4. In fact when equation 9 evolves over time, the level set function ϕ will converge to $-\infty$ inside the curve and to $+\infty$ outside. To avoid dealing with too large values, ϕ can be truncated, for example restricted to $[-100, 100]$. Althought, one must be careful as we remarked that if ϕ is truncated to a small value (eg. to $[-1, 1]$), it may lead to numerical artifacts. It is also possible to use a semi-implicit time discretization, which also increases the stability [26]. In this scheme, the value $\phi_{i,j}$ of the central pixel is taken at time $n + 1$ for the right hand side of equation 7. This yields

$$\phi_{i,j}^{n+1} = \phi_{i,j} + \frac{\Delta t |\nabla\phi|_{i,j} [d_1 \phi_{i+1,j}^n + d_2 \phi_{i-1,j}^n + d_3 \phi_{i,j+1}^n + d_4 \phi_{i,j-1}^n + ((g - c_1)^2 - (g - c_2)^2)]}{1 + \Delta t |\nabla\phi| (d_1 + d_2 + d_3 + d_4)}.
 \tag{10}$$

We will analyze the convergence of equation 9 to a steady state using the semi-implicit time discretization described in equation 10 under several factors. We will study the influence of the initial condition and of the fidelity term chosen. We will also propose an iterative method to obtain multi-phase segmentation.

In [27], Song and Chan decoupled equation 9 in order to simplify the resolution. They first smooth the image g to obtain a new image g^* , and then solve the simplified equation

$$\phi_t = (g^* - c_1)^2 - (g^* - c_2)^2.
 \tag{11}$$

They proposed a fast algorithm for solving equation 11. In fact, their algorithm corresponds closely to clustering the set $\{g^*(x) : x \in \Omega\}$ of values of g^* into two

clusters using a k-means procedure [15]. Osher and He introduced a similar algorithm to solve the multi-phase problem [16]. Again, this is closely related to applying a k-means procedure to the set $\{g^*(x) : x \in \Omega\}$ with the right number of clusters k . Independently, Gibou and Fedkiw took advantage of this analogy to propose an hybrid algorithm that alternates between smoothing and k-means [13].

However, it is not known if any of these fast algorithms lead to minimums of the Chan-Vese energy for some value of the parameters μ and λ .

2. IMPROVEMENTS

We focus on the resolution of the Chan-Vese model via equation 9 in order to extract the heart shape from a 3D CT scan of a human thorax. We propose improvements on several aspects of the resolution process in order to fit the needs of large 3D applications. These aspects are:

- (1) Iterative segmentation instead of multi-phase segmentation.
- (2) Choice of a fidelity term in equation 9.
- (3) Choice of an initial condition in equation 9.

2.1. Iterative Segmentation. Medical images usually contain many regions of different pixel intensity value. In this case a two-phase segmentation will generally not be able to extract the region of interest.

A solution to the multi-region problem is the algorithm proposed by Chan and Vese for multi-phase segmentation that simultaneously evolves several level set functions, that is several curves [28]. There are drawbacks to this method. One is that instability may appear from solving these PDE's simultaneously. Also it may happen that the two curves need to coincide at some places. In this case, if the curvature term is dominant, it may lead to miss-classification in this region, since the curves may not superimpose correctly. Figure 1 shows an example of this phenomena.

The other approach is to look at several level curves of the level set function ϕ [10]. If n different values are chosen, it splits the domain into $n+1$ different regions. One possible issue is that it can not recover triple junctions, as level curve are naturally parallel to each other.

Instead of these strategies, we will segment by solving the simple two-phase Chan-Vese problem iteratively. The process is as follows:

- (1) Split the domain Ω into Ω^+ and Ω^- using the 2-phase Chan-Vese model.
- (2) Stop if the object is extracted, that is if the object is either Ω^+ or Ω^- . Otherwise, decide which of Ω^+ or Ω^- contains the object of interest, and pick this region as a new domain for step 1.

Figure 2 shows the results of an iterative segmentation process on a slice of a CT scan of the heart.

There are many advantages to proceed in this manner. First, if we are interested in a single object in the image, the computations need not to be done on the whole domain after the first step. This clearly saves on CPU time. We can thus focus on the region of interest, rejecting at each step the part of the image that does not contain the given object.




FIGURE 1. An example of an image for which segmentation with two level set function may lead to incorrect segmentation. (a) is the synthetic image to be segmented. (b) shows a correct segmentation of the image, when the curvature term is not dominant ($\mu = \lambda = 1$). (c) is the result if the curvature term is more important ($\mu = 10000$, $\lambda = 1$): there is a miss-classification of some pixels. (d) shows a close-up on the curves where there is miss-classification: it comes from the fact that level curves are only nearly superposed in these regions.

However a full 2^n -phase segmentation of the image can also be done with this iterative method. The first step splits the domain into two different sub-domains Ω^+ and Ω^- . In the second step, each of these sub-domains will be split in 2 parts, yielding four sub-domains Ω^{++} , Ω^{+-} , Ω^{-+} and Ω^{--} . This can be done until the requested number of sub-domains is reached. It takes n such steps to get 2^n regions. In the multi-phase case, n different level sets functions need to be evolved. One step of the iterative method has the same computational cost as evolving a single level set function in the multi-phase method. However in practice the multi-phase




FIGURE 2. 2D Iterative segmentation of the heart from a CT scan image. (a) The image to be segmented (b) the result of the first application of the 2-phase Chan-Vese model, the blue part is the region of interest (c) Second step: in green is the region that has been ignored in the segmentation. The red region will be chosen. (d) Red is chosen (e) blue is chosen (f) the blue region is the heart muscle, we stop.

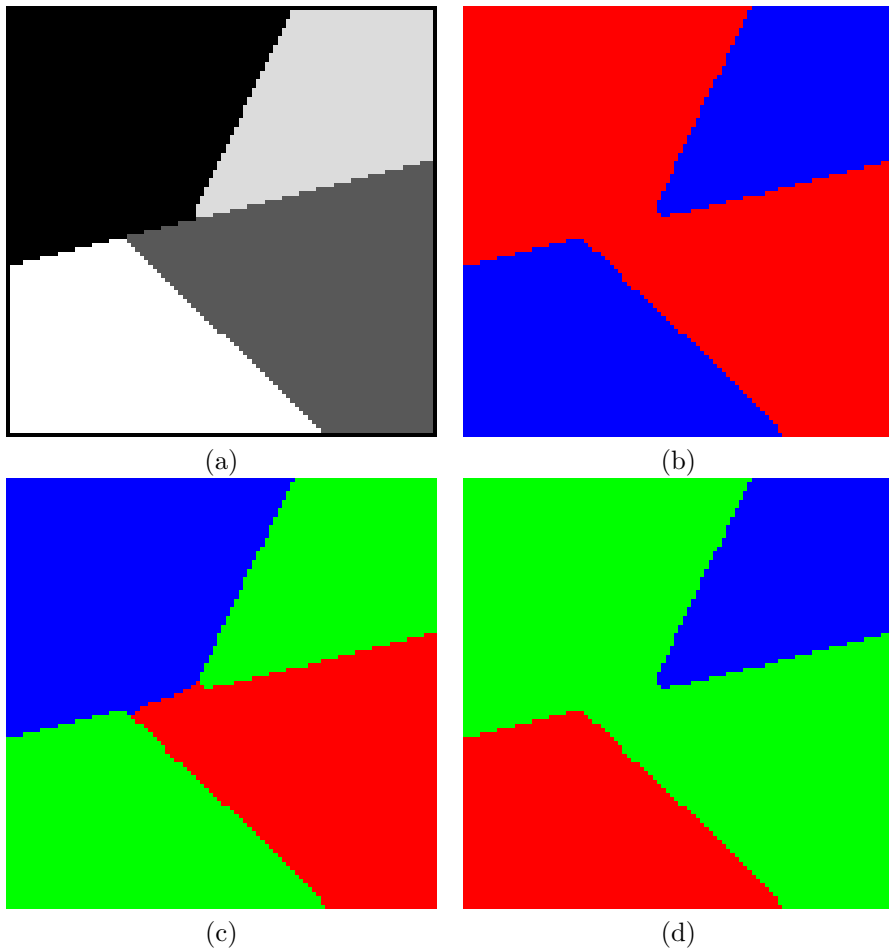


FIGURE 3. Iterative segmentation of the synthetic image (a). (b) shows the two-phase Chan-Vese model applied to the image ($\mu = 10000$, $\lambda = 1$). (c) is the result of the two-phase algorithm applied to the red region of (b) ($\mu = 100000$, $\lambda = 1$) and (d) shows the two-phase segmentation of the blue part of (b) ($\mu = 10000$, $\lambda = 1$). Note that the parameters are different, since they mainly depend on the size of the fidelity term.

algorithm requires smaller time step to ensure stability, which makes it slower. Another advantage of the iterative method is that the smoothing parameters can be varied from steps to steps. This saves time since the equation can be solved with larger time steps when the smoothing term is not dominant.

Another benefit of our iterative method is that triple junctions are accurately obtained since the boundary conditions force the curve to be normal to the domains boundary. Figure 3 shows how our method performs on the synthetic image of Figure 1. It is clear that the miss-classification problem is avoided with the iterative method.

2.2. L^1 fidelity term. Replacing the classical L^2 fidelity term by an L^1 fidelity term is an idea that has first been introduced in signal processing [1, 2, 3] and later in image processing [21, 22]. Using L^1 fidelity term makes the problem more robust to noise and outliers in the signal or image.

It has been remarked that for TV denoising, a L^1 fidelity is more natural [8], since in this case the problem is scale-invariant. Recently, it has been shown that the Chan-Vese model benefits from the same properties when L^2 fidelity is replaced by L^1 fidelity. Indeed if (ϕ, c_1, c_2) minimizes

$$(12) \quad E_{CV,g}^{L^1}(\phi, c_1, c_2) = \int_{\Omega} |g - c_1| H(\phi) dx + \int_{\Omega} |g - c_2| (1 - H(\phi)) dx + \int_{\Omega} |DH(\phi)| dx,$$

then $(\phi, \lambda c_1, \lambda c_2)$ minimizes $E_{CV,\lambda g}^{L^1}$ [11]. This seems to be a desirable property since the scaling of the image should not affect the geometry of the segmented objects.

Computing the Euler-Lagrange equation of 12 yields to the gradient descent equation

$$(13) \quad \phi_t = |\nabla\phi| \left[\operatorname{div} \left(\frac{\nabla\phi}{|\nabla\phi|} \right) + |g - c_1| - |g - c_2| \right],$$

where $\delta(\phi)$ has again been replaced by $|\nabla\phi|$. Now c_1 and c_2 are the median values of g in $\{\phi \leq 0\}$ and $\{\phi > 0\}$ respectively instead of the mean values. These values are to be updated at each time step. Computing the median is more demanding than the mean, but this is negligible compared to the time required to compute the update of the level set function at each iteration.

How does one decide which fidelity term to chose? The L^2 -fidelity is more sensible to noise as it can be seen in Figure 4. It comes from the fact that the mean value is affected by large values whereas the median is affected by frequent values. Hence, if the object to be segmented has a very distinct color, the L^2 -fidelity term should be preferred. If the object has color close to the rest of the image or if the image is very noisy, the L^1 -fidelity will be more efficient. The CT scan of the thorax falls in the second category. We illustrate this by doing an iterative segmentation of a 2D slice using the L^1 -fidelity, see Figure 5. Note that the number of steps needed for the heart segmentation drops from 5 to 3.

2.3. Initial condition. The Chan-Vese energy may admit some local minimums, depending on the image g and on the scale of the parameter μ . The choice of a good initial condition is then crucial in order for the curve not to get stuck into a local minimum which is not the global minimum. The convergence to the steady state may also be slowed down by a bad choice of initial condition. We tried four different strategies and compared the results obtained. They are summarized in figure 6.

The first strategy consists of taking a circle as an initial curve and its signed distance function as level set function. Many level set functions may describe the curve, but the signed distance function has the advantage of being the most regular.



FIGURE 4. Segmentation of the cameraman image using $L2$ -fidelity (left) and $L1$ -fidelity (right), with the same weights ($\mu = 1$ and $\lambda = 0.1$)

The second strategy is to take the union of many circles spread out over the image as an initial curve. Again the signed distance function to this curve is taken as level set function. The idea is that if the curve is spread out, it should be close to features of interest. The first two strategies are used in the seminal paper of Chan and Vese [9].

In the same direction, we propose to take as initial condition a level set function that takes random values in $[-1, 1]$ over the domain Ω . Then there are point inside and outside the curve almost everywhere. The level set function is not regular, but this doesn't seem to be a problem since the level curves gets smooth very fast as time evolves.

The last strategy comes from the fact that when $\mu = 0$, the Chan-Vese problem is very easy to solve via a k-means procedure. This splits Ω into two regions Ω_1 and Ω_2 . One can take as an initial condition $\phi = \chi_{\Omega_1} - \chi_{\Omega_2}$. The signed distance function to $\partial\Omega_1 \cap \partial\Omega_2$ would be another option, but numerical experiments suggest that it is a better idea that all points are initially relatively close to 0.

We tried the different strategies on the well known cameraman picture. On Figures 7 and 8 we compared the decay of the energy and the L^2 -convergence to the absolute minimum. The energy is approximated using a discretization of the equation 3, the term $DH(\phi)$ is simply discretized using a forward Euler scheme. If a regularized version $H_\epsilon(\phi)$ of $H(\phi)$ is used, the energy becomes less accurate as the function ϕ evolves since ϕ develops sharp gradients. Figure 9 shows the semi-log plot of the L^2 distance to the absolute minimizer.

On these figures, it is possible to see two different behaviors. At the beginning of the iteration process the fidelity term is dominant in equation 9. In this case, the random initial curve and the ($\mu = 0$) solution converge very quickly towards the steady state, as seen in Figure 8. As the curves get closer to the steady state,




FIGURE 5. Iterative segmentation of the heart from a 2D slice of a CT scan using the L^1 -fidelity. 3 steps are required instead of 5 for the L^2 -fidelity.

the curvature term gain importance and then every curve approaches the steady state at a similar rate as shown on the semi-log plot in Figure 9. The same kind of phenomena happens for the energy.

It is not easy to determine which initial condition is better than the other. It seems that the random curve and the $(\mu = 0)$ solution converge quickly to a state that is close to the steady state, which is certainly a nice property. However in practice we favor the use of the random curve since it has no *a priori* on the position of the interface and the points inside and outside the curve are evenly spread over the image. In this case, there is a lower chance to get stuck in a local minimum.

3. 2D APPLICATION

As an illustration of the method, we have presented in Figure 5 the results of the iterative segmentation with L^1 norm on a 2D CT slice. In order to use the




FIGURE 6. The different initial curves: (a) one circle, (b) 100 equidistant circles, (c) random values in $[-1, 1]$ and (d) the solution of the problem when $\mu = 0$ and $\lambda = 1$.

results of the segmentation process as a base for numerical simulations of the organ functions, it is usually required to generate a mesh of the given organ. To do so, we used the simple and very efficient 2D/3D mesh generator DistMesh [25, 17].

Distmesh is a Matlab tool designed for meshing domains implicitly defined through a level set function. Most mesh generators start by meshing the boundary of the domain (a curve for 2D domains, a surface for 3D domains) and then march to mesh the interior. DistMesh has a different approach: it triangulates the whole domain and then moves vertices to well fit the boundary, with a control on the size and quality of the elements. It can also be modified to accommodate sub-domains. For each sub-domain, we have computed the signed distance function to the boundary. The signed distance function is computed using the reinitialization equation

$$(14) \quad \phi_t = S(\phi)(1 - |\nabla\phi|),$$

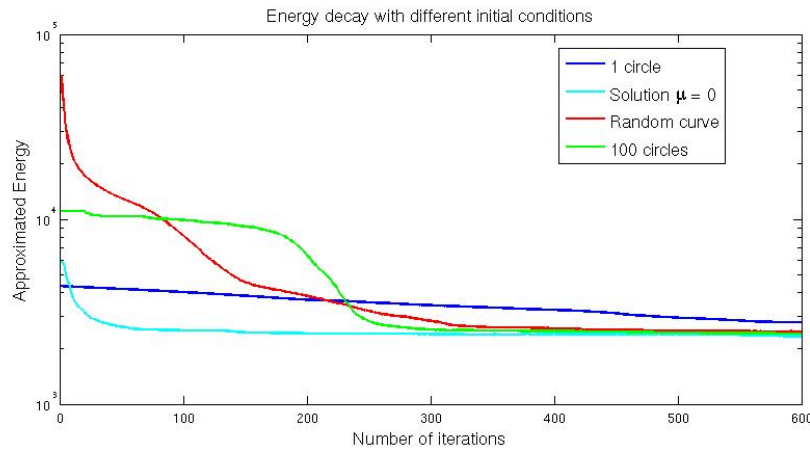
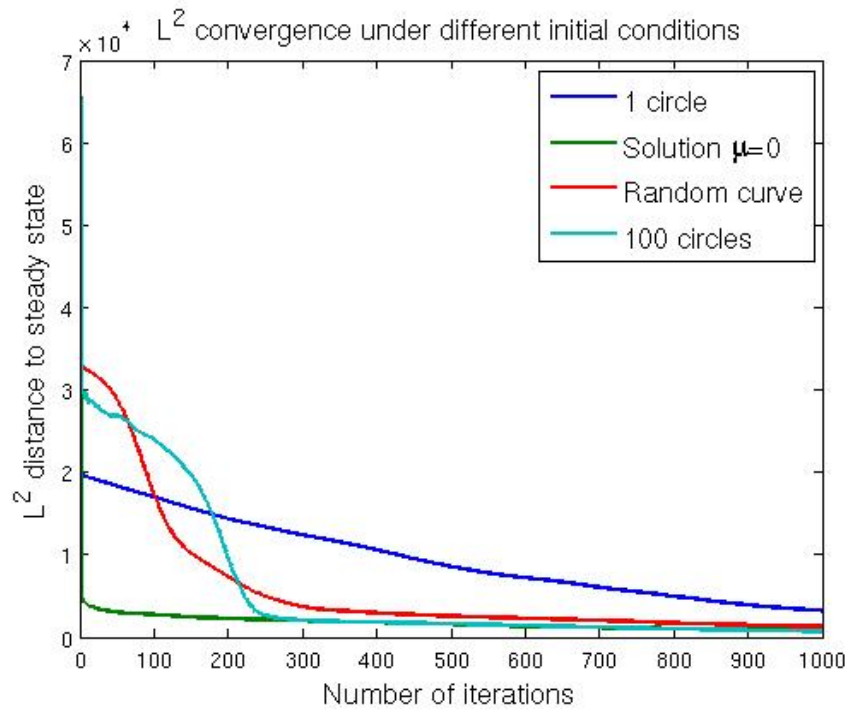


FIGURE 7. The energy decay under various initial conditions.

FIGURE 8. The L^2 -convergence under various initial conditions.

where $S(\phi)$ is a regularization of the sign function [23]. This signed distance function is used to project on the boundary nearby vertices. It is then sufficient to compute the signed distance function only in a neighborhood of the zero level set.

With DistMesh, an element diameter has to be specified. This one may differ spatially. We remarked that nice boundaries can be obtain if the element size is

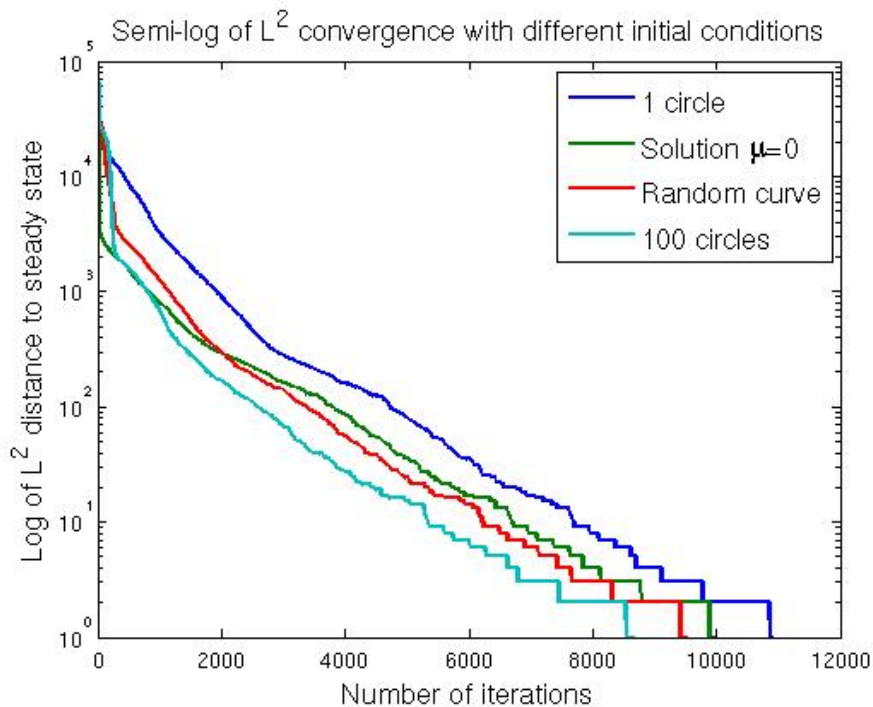



FIGURE 9. The semi-log plot of the L^2 -convergence under various initial conditions.

decreased continuously when approaching sub-domains boundaries. In practice, the element size is usually decreased to half the size it has in the sub-domain. Figure 10 shows an example of a 2D mesh of the thorax, built from the previous segmentation in Figure 5.


4. 3D APPLICATION

We will apply our framework to 3D segmentation of a full CT scan. This CT scan is by courtesy of the Heart Institute of the University of Ottawa. It is of size $512 \times 512 \times 199$ and the voxels have a 0.48mm resolution in the transverse plane (the $x - y$ plane) and a 1.25mm. resolution in the direction of the transverse axis (z direction). This image has over 52 million voxels, it is thus a challenge to solve the segmentation problem in an efficient way.


However, since the equation is solved with an explicit time scheme, it is possible to parallelized the solver. The idea is to split the image into pieces that are to be distributed on computing nodes. A simple choice is to split the z direction into as many blocks as we want. At each iteration, the computations can be made independently on each block. Then only the information about shared boundaries of the block need to be exchanged. In the case where the image is split in only one direction, there is only the information about the top slice and the bottom slice to be exchanged, which can be done very quickly.



(a)



(b)



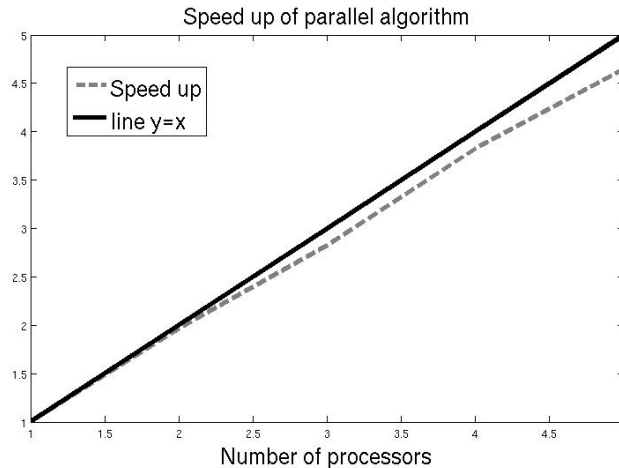


FIGURE 11. The speed up for the parallel algorithm: the time for doing an iteration with one processor divided by the time for doing an iteration with x processors. It is just below the perfect speedup, that is the line $y = x$.

If we apply the Chan-Vese model with L^2 fidelity, we need to compute the average of a function in a given region. This is easy to do in parallel. However, if we are to use the L^1 fidelity, we need to compute the median in parallel, which is trickier. To find the median in the serial case, the C++ template `nth_element` is used [18], which is a linear algorithm for partial sort (it is more efficient than complete sort which is of order $n \log(n)$, where n is the number of pixels). Hence for preliminary tests, we used the L^2 fidelity. For all computations, a random initial curve has been chosen. The computations then require less time steps.

The parallel code is implemented in C++ using the openMPI library [14]. The amount of RAM memory required to solve the Chan-Vese problem is then divided into the different nodes. The speed up for this parallelization is nearly perfect: using n processors divides the time per iteration by n , see figure 11.

Most computations are made on a 16 processor SUN cluster with distributed memory. Running the code on 6 processors divides the CPU time by almost 6. As the number of processors increases, the time lost in data transmission becomes more important. Note that many computations have been made successfully on a simple dual-core laptop, cutting the CPU time in 2, although the same amount of RAM is required.

Figure 12 shows the 3D CT image as well as the result of the first segmentation step. Following this first segmentation, one side of the surface must be chosen as the new segmentation domain. The interior is the region that contains the trachea, the exterior contains the heart. To obtain the trachea we segmented the interior region with a high curvature term, since there is much noise and irregularities in the lunges. Once the segmentation process is complete, we just need to choose the connected component that corresponds to the trachea. Figure 13 shows the




FIGURE 12. The CT scan to be segmented (a) and the result of the first segmentation step (b)

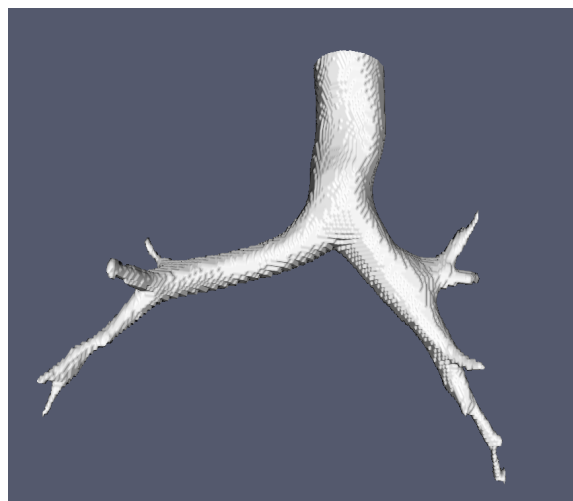


FIGURE 13. The trachea segmented from the 3D CT scan.

resulting segmentation. The trachea is then meshed with DistMesh. Figure 14 shows 2 different meshes of the trachea: a coarse one of about 40 000 tetrahedra, and a finer one of about 500 000 tetrahedra.

Now, if we want to segment the heart, we have to take the result of the first segmentation and choose the exterior of the surface. To get to the heart, four more segmentation steps will be required (five steps in total). If this would be done with several level set functions, it would then take 5 level set functions to get this level of details, which is hard to solve on images of this size.




FIGURE 14. Two meshes of the trachea generated with DistMesh: a coarse mesh of about 40 000 tetrahedra (a) and a finer of about 500 000 tetrahedra.




FIGURE 15. 2 views of the exterior surface of the heart in the final segmentation. The general shape is well recovered.

From the second stage, we have decided to do the segmentation not on the exact image but on a blurred version of it, due to the high level of noise. To blur the image, we apply the linear heat equation on the image for a few time steps with an explicit scheme. This is also easily done in parallel. Figure 15 shows the exterior surface of the heart in the final segmentation and Figure 16 shows the segmented heart cavities. The cavities are well segmented, especially the one of left ventricle, which is the most important part in many applications. The position of the mitral valve is also precisely captured as shown in Figure 17. The general shape of the heart is well captured, ventricles and atria are extracted as well as the aorta.

There are some imperfections in the segmentation of the heart surface near the epicardium, as the surface leaks in some areas. This is due to the fact that the heart touches the liver at that place in the image and that the two organs are of same grey level. Further work has to be done to fully extract the heart.




FIGURE 16. 2 views of the interior surfaces of the heart in the final segmentation. The pillars in the left ventricle are well segmented.




FIGURE 17. Two other views of the internal surface of the final segmentation: a view of the position of the mitral valve (a) and a combined view of the left ventricle and the left atria (b).

CONCLUSION

In his paper, the efficiency of the iterative version of the Chan-Vese model have been shown on test cases and successfully applied to large 2D and 3D applications to the human trachea and heart. We also conclude that the L^1 fidelity is more efficient for this kind of segmentation. It is less sensitive to noise and to small regions of different color intensity. Analysis of convergence under various initial conditions showed that our new approach which consists of taking a random initial curve is very efficient. The curve quickly converges to a state close to the steady

state and has no shape prior. It has been successfully apply in 2D and 3D for CT scan segmentation. We also explained how the algorithm is implemented in parallel and the computing gain obtained this way. Finally, high quality fine 2D and 3D meshes have been produced. These meshes are actually used for Finite Element and Finite Volume simulations in electro-physiology.

ACKNOWLEDGMENT

The authors would like to thank the Ottawa Heart Institute for providing them with quality CT scans of the thorax. The authors acknowledge the support of NSERC under a Discovery Grant and a CGSD (Canada Graduate Scholarship for Doctoral studies) post graduate scholarship for the first author.

REFERENCES

- [1] S. Alliney. Digital filters as absolute norm regularizers. *Signal Processing, IEEE Trans. Signal Process.*, 40(6):1548–1562, 1992.
- [2] S. Alliney. Recursive median filters of increasing order: a variational approach. *Signal Processing, IEEE Trans. Signal Process.*, 44(6):1346–1354, 1996.
- [3] S. Alliney. A property of the minimum vectors of a regularizing functional defined by means of the absolute norm. *IEEE Trans. Signal Process.*, 45(4):913–917, 1997.
- [4] L. Ambrosio, N. Fusco, and D. Pallara. *Functions of Bounded Variation and Free Discontinuity Problems*. Oxford University Press, USA, 2000.
- [5] L. Ambrosio and V.M. Tortorelli. Approximation of functionals depending on jumps by elliptic functionals via Γ -convergence. *Comm. Pure Appl. Math*, 43(8):999–1036, 1990.
- [6] L. Ambrosio and VM Tortorelli. On the approximation of free discontinuity problems. *Boll. Un. Mat. Ital*, 6:105–123, 1992.
- [7] A. Braides. *Approximation of Free-Discontinuity Problems*. Springer, 1998.
- [8] T.F. Chan and S. Esedoglu. Aspects of total variation regularized L1 function approximation. *SIAM J. Appl. Math*, 65(5):1817–1837, 2005.
- [9] TF Chan and LA Vese. Active contours without edges. *IEEE Transactions on Image Processing*, 10(2):266–277, 2001.
- [10] G. Chung and L.A. Vese. Energy minimization based segmentation and denoising using a multilayer level set approach. *Lecture Notes in Computer Science*, 3757:439–455, 2005.
- [11] J. Darbon. A Note on the Discrete Binary Mumford-Shah Model. *Lecture Notes in Computer Science*, 4418:283, 2007.
- [12] O. Ecabert, J. Peters, C. Lorenz, J. von Berg, M. Vembar, K. Subramanyan, G. Lavi, and J. Weese. Towards automatic full heart segmentation in computed-tomography images. *Computers in Cardiology, 2005*, pages 223–226, 2005.
- [13] F. Gibou and R. Fedkiw. A fast hybrid k-means level set algorithm for segmentation. *4th Annual Hawaii International Conference on Statistics and Mathematics*, pages 281–291, 2005.
- [14] W. Gropp, E. Lusk, and A. Skjellum. *Using MPI: Portable Parallel Programming with the Message Passing Interface*. MIT Press, 1999.
- [15] J.A. Hartigan. *Clustering Algorithms*. John Wiley & Sons, Inc. New York, NY, USA, 1975.
- [16] L. He and S. Osher. Solving the Chan-Vese Model by a Multiphase Level Set Algorithm Based on the Topological Derivative. *Proceedings of the 1st International Conference on Scale Space Variational Methods in Computer Vision*, 2007.
- [17] <http://www-math.mit.edu/~persson/mesh/>.
- [18] http://www.sgi.com/tech/stl/nth_element.html.
- [19] A. Marquina and S. Osher. Explicit algorithms for a new time dependent model based on level set motion for nonlinear deblurring and noise removal. *SIAM J. Sci. Comput.*, 22(2):387–405, 2001.
- [20] D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Comm. Pure Appl. Math*, 42(5):577–685, 1989.
- [21] M. Nikolova. Minimizers of Cost-Functions Involving Nonsmooth Data-Fidelity Terms. Application to the Processing of Outliers. *SIAM Journal on Numerical Analysis*, 40(3):965–994, 2002.

- [22] M. Nikolova. A Variational Approach to Remove Outliers and Impulse Noise. *Journal of Mathematical Imaging and Vision*, 20(1):99–120, 2004.
- [23] S. Osher and R.P. Fedkiw. *Level Set Methods and Dynamic Implicit Surfaces*. Springer, 2003.
- [24] S. Osher and J. Sethian. Fronts propagating with curvature-dependent speed- Algorithms based on Hamilton-Jacobi formulations. *Journal of Computational Physics*, 79(1):12–49, 1988.
- [25] P.O. Persson and G. Strang. A Simple Mesh Generator in MATLAB. *SIAM Review*, 46(2):329–345, 2004.
- [26] P. Smereka. Semi-Implicit Level Set Methods for Curvature and Surface Diffusion Motion. *Journal of Scientific Computing*, 19(1):439–456, 2003.
- [27] B. Song and T. Chan. A fast algorithm for level set based optimization. *UCLA CAM Report*, 2(68), 2002.
- [28] L.A. Vese and T.F. Chan. A Multiphase Level Set Framework for Image Segmentation Using the Mumford and Shah Model. *International Journal of Computer Vision*, 50(3):271–293, 2002.
- [29] L. Zhukov, Z. Bao, I. Guskov, J. Wood, and D. Breen. Dynamic deformable models for 3D MRI heart segmentation. In *SPIE Medical Imaging*, volume 4684, pages 1398–1405, 2002.