Active Contour Based Segmentation of 3D Surfaces

Matthias Krueger, Patrice Delmas, and Georgy Gimel'farb

Dept. of Computer Science, Tamaki Campus
The University of Auckland, Auckland, New Zealand
mkru007@aucklanduni.ac.nz

Abstract. Algorithms incorporating 3D information have proven to be superior to purely 2D approaches in many areas of computer vision including face biometrics and recognition. Still, the range of methods for feature extraction from 3D surfaces is limited. Very popular in 2D image analysis, active contours have been generalized to curved surfaces only recently. Current implementations require a global surface parametrisation. We show that a balloon force cannot be included properly in existing methods, making them unsuitable for applications with noisy data. To overcome this drawback we propose a new algorithm for evolving geodesic active contours on implicit surfaces. We also introduce a new narrowband scheme which results in linear computational complexity. The performance of our model is illustrated on various real and synthetic 3D surfaces.

1 Introduction

1.1 Motivation

Recently, there has been an increasing interest in generalising methods that are well-known and understood for planar images, e.g. isotropic and anisotropic diffusion [1,2], scale space analysis [2,3] or image inpainting [4], onto images painted on curved surfaces. The goal is to adequately process images from, for instance, 3D scanners or surface reconstruction algorithms.

However, due to the Gauss' famous *theorema egregium* (see e.g. [5]) surfaces with a non-zero Gaussian curvature cannot be mapped onto the plane without metric distortions. Moreover, looking at particular applications, recent research [6,7,8] shows that multi-modal approaches, combining 2D (texture) and 3D (shape) information, perform better than either 2D or 3D techniques alone in face biometrics. Thus, it is necessary to study object segmentation approaches, in particular active contour models, that incorporate both 2D and 3D information.

Object segmentation is one of the foremost problems in computer vision, and Spira and Kimmel [9] first attempted to apply active contours , namely geodesic active contours (GAC), to images on parametric surfaces. While working on parametric surfaces is an intuitive approach, it has a number of drawbacks. In Sect. 4 we will explain in more detail the most serious one: the problems

D. Forsyth, P. Torr, and A. Zisserman (Eds.): ECCV 2008, Part II, LNCS 5303, pp. 350–363, 2008.

arising when a balloon term pushing the contour with a constant force is added. Without a balloon term, however, the application of the method is limited to data involving no or very little noise.

Krueger et al. [10] proposed an equation for GACs on implicitly represented surfaces. Their approach avoids the problems of the parametric approach, but suffers from low computational efficiency. Applying Cheng et al.'s [11] framework for geometric curve evolution on implicit surfaces, Krueger et al. [10] obtain an algorithm with quadratic complexity. For large datasets the algorithm results in processing times which are infeasible in practice (see Sect. 5).

In this article we take up the implicit approach of Krueger et al. [10] and propose a new numerical scheme to reduce its computational complexity to linear. Our efficient implementation in conjunction with the versatility and robustness of the implicit surface representation makes our model applicable to a wide range of synthetic and natural surfaces.

1.2 Related Work

Active contours (also known as "snakes") were introduced in the pioneering work by Kass et al. [12]. Derived from the gradient flow of an energy functional, a snake is a curve that deforms its shape to minimise the total internal and external energy. The internal energy imposes regularity on the curve while the external energy attracts the curve to desired image features like edges. Later, Caselles et al. [13] introduced the GAC, based on an energy functional $\mathcal{E}$ that integrates a positive edge indicator function f along the contour C:

$$\mathcal{E}(C) = \int_C f(C)\, ds\ . \tag{1}$$

This formalisation yields the evolution equation:

$$\partial_t C = (\kappa f - \langle \nabla f, \nu \rangle)\, \nu\ , \tag{2}$$

where $\partial_t C := \frac{\partial C}{\partial t}$, κ, and ν are the time derivative, the curvature, and the unit normal field, respectively. The function f originally had small values for large image gradients and higher values elsewhere. Equation (2) is the gradient descent flow of the energy functional in (1), i.e. a curve evolving according to (2) decreases its energy as fast as possible. It converges towards a steady state if a sufficiently stable minimum of (1) is reached. One major drawback of snakes, their dependency on their parameterisation, was overcome by applying a geometric flow, namely a weighted curvature flow, to regularise the evolving curve. Caselles et al. [13] applied the level sets method [14] allowing the curve to change its topology during the evolution. They further proposed a general GAC model including a balloon force which shrinks or inflates the curve by a constant speed c:

$$\partial_t C = (f \cdot (\kappa - c) - \langle \nabla f, \nu \rangle)\, \nu\ . \tag{3}$$

Spira and Kimmel [9] applied the GAC model (2) to segment images painted on parametric surfaces. The evolution equation in this generalised setting,

$$\partial_t C = (\kappa_g f - \langle \nabla f, \hat{\nu} \rangle)\, \hat{\nu}\ , \tag{4}$$

is similar to (2). Here, κ_g denotes the geodesic curvature and $\hat{\nu}$ is the unit co-normal vector of C on the curved surface M. The co-normal $\hat{\nu}$ is defined as the normalised cross product of the normal ν and the tangent vector field C_s. The geodesic curvature is the tangential component of the curvature. However, it is beyond the scope of this paper to cover the geometric framework of curves on surfaces (see e.g. [9] and references therein). Spira and Kimmel [9] projected (4) onto the parameterisation plane of M and solved the resulting PDE applying the level sets framework [14].

Bogdanova et al. [2] pursued a similar approach. They considered GAC on three particular parametric surfaces, namely, the sphere, hyperboloid, and paraboloid, and applied their results to omnidirectional images. In contrast to Spira and Kimmel [9], Bogdanova et al. [2] derived the flow equations directly from an adapted version of the GAC energy functional (cf. (1)).

Cheng et al. [11] studied basic geometric curve evolution models on implicit surfaces, including the geodesic curvature, advection and normal flows.

Jin et al. [15] studied segmentation on surfaces in a slightly different context. To reconstruct a 3D scene from multiple 2D views in a variational framework, they simultaneously evolved both a level set surface and a radiance pattern of the surface. Discontinuities of the radiance were modeled by curves on the surface. Jin et al. [15] applied Cheng et al.'s [11] framework to evolve curves on the surface. However, their model is only 2.5D because the actual image segmentation according to a Mumford-Shah model takes place after projecting the surface to the original 2D views. Their approach is not applicable to general images on implicit surfaces since in general 2D projections are not available.

Krueger et al. applied Cheng et al.'s [11] work to propose a GAC model for images on implicit surfaces as succinctly outlined in [10]. We will review this approach in more detail in Sect. 2.

1.3 Contribution of the Paper

In this paper we present a new framework for feature extraction from 3D surface data using GACs. Working with implicitly defined surfaces we avoid numerical difficulties inherent to the parametric approach. Most importantly, in contrast to Spira's work [9] a balloon force can be integrated easily in our approach, making it applicable to real data which are usually noisy. To reduce the computational complexity of the implicit approach to linear we propose a new narrowband algorithm. Our experimental results show that in contrast to previous works our model can be applied to most of the complex surfaces in practice. To underline that, we present results of applying our 3D GAC model to human faces as well as other intricate, non-synthetic surfaces. To the best of our knowledge, this has not been done before.

The paper is organized as follows. Section 2 reviews Cheng et al.'s [11] framework for geometric curve evolution on implicit surfaces and the GAC model proposed by Krueger et al. [10]. Section 3 gives details on the numerical implementation and introduces the new narrowband scheme. Theoretical considerations comparing our approach with the known works on parametric surfaces are

made in Sect. 4. Section 5 presents results of our experiments on synthetic and natural surfaces, and Sect. 6 concludes the paper.

2 Geodesic Active Contours on Implicit Surfaces

After embedding a curve C as an isocontour in a 2D scalar function ϕ, the evolution of the curve is controlled by solving a suitable PDE for ϕ (see e.g. [14,16]). The GAC in (3) corresponds to the following level set PDE:

$$\partial_t \phi = \underbrace{cf|\nabla\phi|}_{I} + \underbrace{\langle \nabla f, \nabla\phi \rangle}_{II} + \underbrace{f\kappa|\nabla\phi|}_{III} , \tag{5}$$

where the terms I, II, and III denote a motion in normal direction, an advection flow, and a weighted mean curvature flow, respectively. In order to generalise (5) to implicit surfaces, Krueger et al. [10] applied Cheng et al.'s framework [11] for geometric curve evolution on implicit surfaces. The latter forming the backbone of the model discussed in this paper is briefly reviewed below.

Given an implicit surface $M = \{x \in \mathbb{R}^3 | \psi(x) = 0\}$, an important question is how to describe a curve on M. Since there is no parameterisation at hand, the level set method has to be extended to this case. Cheng et al. [11] proposed to describe a curve on an implicit surface by intersecting the surface with another implicit surface given by the level set function ϕ. This is shown for two example surfaces in Fig. 1.

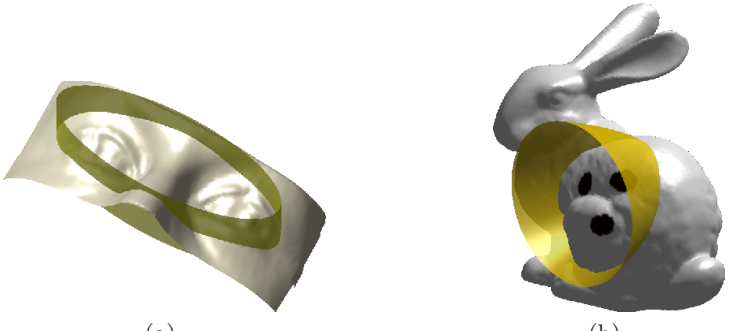

(a) (b)

Fig. 1. Initialisation of a curve on an implicitly given face (a) and bunny (b) by intersecting the surface with a cylinder and a paraboloid, respectively

It is well known [14,16] that the curvature flow equation in the level set setting is

$$\partial_t \phi = \left[\nabla \cdot \left(\frac{\nabla\phi}{|\nabla\phi|} \right) \right] |\nabla\phi| , \tag{6}$$

with the term in the square brackets being the curvature. In order to adapt this equation to the surface setting, Cheng et al. [11] projected all differential

operators involved onto the surface M using the projection operator $P_w(v) = v - \frac{1}{|w|^2}\langle v, w\rangle w$, where v and w are vectors in $\mathbb{R}^3$. With $\nabla\psi(x)$ being orthogonal to M at point x, the operator $P_{\nabla\psi(x)}$ projects a vector onto the tangent plane T_xM of M at x. Cheng et al. [11] also showed that for a real valued function u defined on $\mathbb{R}^3$, the surface gradient ∇_M is given by $\nabla_M u = P_{\nabla\psi}\nabla u$. The surface divergence can be handled analogously. Projecting equation (6) onto the surface M yields the following formula for the geodesic curvature flow:

$$\partial_t\phi = \underbrace{\left[\nabla\cdot\left(\frac{P_{\nabla\psi}\nabla\phi}{|P_{\nabla\psi}\nabla\phi|}|\nabla\psi|\right)\right]\frac{1}{|\nabla\psi|}}_{\kappa_g}|P_{\nabla\psi}\nabla\phi| \;, \tag{7}$$

where κ_g is the geodesic curvature. Similar considerations lead Cheng et al. [11] to the equations for the geodesic advection flow under a static vector field V and the constant normal flow with velocity V_0. Those are given by $\partial_t\phi = \langle P_{\nabla\psi}V, \nabla\phi\rangle$ and $\partial_t\phi = V_0|P_{\nabla\psi}\nabla\phi|$, respectively.

Applying Cheng et al.'s [11] framework, Krueger et al. [10] generalized the GAC model (5) to implicit surfaces. With a surface M embedded in a cube $\Omega \subset \mathbb{R}^3$, a greyscale image I is a mapping $I : M \to \mathbb{Z}^+$. To apply Cheng et al.'s [11] method, the image I has to be smoothly extended to at least a certain band surrounding M (see e.g. [16]). As in the classical GAC model, the image I should be smoothed before calculating an edge indicator function. If a textured image is mapped onto the surface, this can be achieved by applying a low pass 2D filter to the image. Alternatively, images can be smoothed directly on an implicit surface [1] or by applying a 3D filter to the volumetric image I.

If I_σ is the smoothed image, the edge indicator function can be defined as $f = \frac{1}{1+|\nabla I_\sigma|^p}, \quad p \in \{1, 2\}$. Note, that f can be defined in terms of geometric quantities such as the surface curvature as well (see our experiments in Sect. 3.1). Then the equation for GACs on implicit surfaces is

$$\partial_t\phi = \underbrace{cf|P_{\nabla\psi}\nabla\phi| + \langle P_{\nabla\psi}\nabla f, \nabla\phi\rangle}_{I} + \underbrace{\mu f\kappa_g|P_{\nabla\psi}\nabla\phi|}_{II} \;. \tag{8}$$

where the parameters c and μ determine the influence of the balloon force and the regularity terms, respectively. The image gradient term has a normalised weight. Then the evolving contour on M is given by the zero level set $C(t) = \{x \in \Omega | \phi(x, t) = \psi(x) = 0\}$. In Sect. 3 we study the numerical implementation of (8) in more detail.

3 Numerical Implementation

The general numerical framework for implementing evolution equations on implicitly given surfaces was introduced succinctly by Cheng et al. [11]. Krueger et al. [10] illustrated the numerical procedures for implementing (8) in more detail. State-of-the-art techniques for solving level set equations [16] were used to obtain a stable numerical algorithm.

Analysing (8), we note the term II is diffusive and can be approximated by central differences. However, term I is of Hamilton-Jacobi (HJ) type, thus the derivatives have to be computed using Essentially Non Oscillating (ENO) or Weighted Essentially Non Oscillating (WENO) schemes to avoid numerical instabilities [14,16]. Following Krueger et al. [10], we approximate the hyperbolic term I by a Lax-Friedrichs (LF) scheme with global flux and the time derivative by forward differencing.

Cheng et al. [11] proposed to reinitialise the level set function in two stages:

1. Transforming ϕ to a signed distance function on M (see (9)),
2. Enforcing ϕ to have its level sets perpendicular to M (see (10)).

Both stages were achieved by iterating a few steps of the adequate partial differential equation, namely

$$\partial_t \phi = \operatorname{sign}(\phi_0)(1 - |P_{\nabla\psi} \nabla\phi|) \tag{9}$$

and

$$\partial_t \phi + \operatorname{sign}(\psi) \left\langle \frac{\nabla\psi}{|\nabla\psi|}, \nabla\phi \right\rangle = 0 \,. \tag{10}$$

In our experiments we observed that fifth-order WENO schemes are necessary to accurately reinitialize the level set function. For the GAC equation (8) second-order ENO proved to be sufficiently accurate.

3.1 The Optimised Framework

Cheng et al. [11] localised the numerical calculations to a neighbourhood of the surface, obtaining quadratic complexity. In this paper, we propose a new algorithm with linear complexity for evolving GACs on implicit surfaces. Essentially, it extends the approaches introduced in [10,11,17].

Let $\widetilde{\psi}$ be a signed distance function for M, which can be computed efficiently by the fast marching method [18]. We also assume ϕ is close to a signed distance function on M at a time t_0, e.g. $t_0 = 0$. We retain $\widetilde{\phi} := \phi(., t_0)$ (see below). The main idea is to define a function θ approximating the distance of a grid point in space from the evolving curve at time t_0. We define θ to be the Euclidean norm of $\widetilde{\psi}$ and $\widetilde{\phi}$, i.e. $\theta(x) := \sqrt{|\widetilde{\psi}(x)|^2 + |\widetilde{\phi}(x)|^2}$. Since both $\widetilde{\psi}$ and $\widetilde{\phi}$ are close to the signed distance functions in $\mathbb{R}^3$ and on the surface M, respectively, the level sets of θ resemble circular tubes along the curve C.

Having defined θ, we localise the level set equations as in [17]: Let $0 < \beta < \gamma$ be two constants which determine the radii of the inner and the outer narrowband. For our setting (with unit grid size $\Delta x = 1$ and fifth-order schemes), we obtained good results with $\beta = 2$ and $\gamma = 4$ (see Fig. 2). Only intricate surfaces with a comparably low grid resolution – such as the Stanford bunny in our experiment in Sect. 5 – required a larger narrowband with $\beta = 3$ and $\gamma = 6$ as suggested in [17] for the use of fifth order WENO schemes.

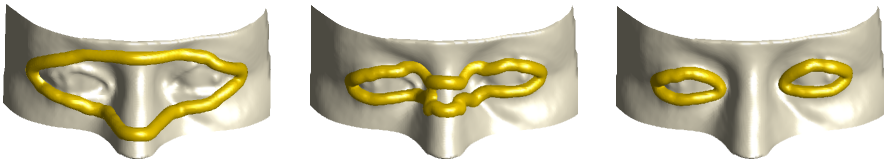

(a) (b) (c)

Fig. 2. Evolving narrowband on a face surface

Our algorithm for evolving GAC on implicit surfaces consists of the following steps:

Step 0. *Initialise distance functions* $\widetilde{\psi}$ *and* $\widetilde{\phi}$.
Since M is static, this has to be done only once for $\widetilde{\psi}$. However, the curve evolves and thus $\widetilde{\phi}$ has to be recomputed regularly (see Step 3).

Step 1. *Compute the narrowband* $T = \{x : \theta(x) \leq \gamma\}$ *and the extended narrowband* $N = \{x : \exists \boldsymbol{e} \in E : \theta(x + \boldsymbol{e}) \leq \gamma\}$ *with* E *the canonical basis of* $\mathbb{R}^3$.
θ only depends on $\widetilde{\psi}$ and $\widetilde{\phi}$. The narrowband has to be updated only if ϕ was reinitialised before and thus $\widetilde{\phi}$ has changed (see Step 3).

Step 2. *Update the level set function* ϕ *by solving a modified equation (11) in* T.
In order to avoid boundary stability issues, we localise the equation by multiplying with a cut-off function λ which equals one in the inner narrowband ($\theta \leq \beta$) and smoothly decreases to zero with $\theta \to \gamma$. A particular choice of λ is given in [17]. Then the modified equation is as follows:

$$\partial_t \phi = \lambda(\theta)\Big(cf|P_{\nabla\psi}\nabla\phi| + \langle P_{\nabla\psi}\nabla f, \nabla\phi\rangle + \mu f \kappa_g |P_{\nabla\psi}\nabla\phi|\Big) \; . \qquad (11)$$

Step 3. *Reinitialise if necessary, then return to Step 1.*
The level set function ϕ is reinitialised iff one of the three criteria is met:

1. The evolving curve reaches the tube $\{x : \theta(x) = 1\}$ of radius 1,
   (this ensures the updated narrowband T does not outreach the current extended narrowband N)
2. $|P_{\nabla\psi}\nabla\phi|$, differs from 1 significantly , or
3. The maximum number of steps is reached.

Reinitialisation is achieved by iterating a few steps of Eqs. (9) and (10) on the extended narrowband N. Since both these equations are hyperbolic, no boundary conditions are needed. Afterwards, we retain $\widetilde{\phi} = \phi(., t)$ and return to Step 1.

4 Implicit vs. Parametric Surfaces

Below we illustrate, that GAC in the parametric framework introduced by Spira et al. [9] have a serious intrinsic drawback.

Previous work on GAC on non-Euclidean geometries using parametric surface representations [9,2] start from adapted versions of functional (1) to derive

generalized equations for the GAC such as (4). Note, that no balloon force is included in this equation. In fact, an analysis of the respective equations shows that incorporating a balloon force in the parametric setting is problematic and results in a non-uniformly scaling equation. Without a balloon force however, GACs mostly fail when data are noisy, unless the contour is initialised very close to the desired object. In Sect. 5 we will support the theoretical considerations of this section with experimental results.

Starting from Spira et al.'s framework [9], for the sake of clarity we assume that the metric tensor g_{ij} is conformally equivalent to the Euclidean metric tensor, i.e. $g_{ij} = k \cdot \delta_{ij}$. This corresponds to a simple rescaling of an image. Then the equation for GAC on parametric surfaces reduces to

$$\partial_t \phi = f \frac{\kappa}{k} |\nabla \phi| + \frac{1}{k} \langle \nabla f, \nabla \phi \rangle \ , \tag{12}$$

where κ is the curvature of the evolving curve in the parameterisation plane. Comparing to the PDE (5) for a standard GAC, the curvature and the advection term scale equally with the factor k^{-1}. In fact, after adaption of the time step, the transformed equation (12) describes the same curve evolution as the original equation ((12) with $k = 1$). The situation is different with the balloon force term: in Spira's [9] framework the equation for the geodesic normal flow on parametric surfaces is given as $\partial_t \phi = -f |\nabla_M \phi|$.[1] After applying $g_{ij} = k \cdot \delta_{ij}$ again, this equation becomes $\partial_t \phi = -k^{-\frac{1}{2}} f |\nabla \phi|$. Note that relative to (5) with $c = 1$, this equation has the scaling factor $k^{-\frac{1}{2}}$. Therefore, adding a balloon term to the GAC equation (12) results in a PDE which does not scale uniformly.

While this is not a problem for standard 2D images with a constant metric, it has undesirable consequences if the metric is variable. In this case the PDE is ill-posed as the ratio between the weights of the balloon and other terms varies. Imagine a model surface with $k = 1$ in certain parts and a bigger $k > 1$ in other parts. Then, the balloon term and the further terms are weighted equally in the area with $k = 1$ (since $k^{-1} = k^{-\frac{1}{2}}$ for $k = 1$), while the balloon term is weighted higher in the area with $k > 1$ (as $k^{-\frac{1}{2}} > k^{-1}$ for $k > 1$). In the next section we will illustrate how this effect causes undesirable segmentation results.

The above problem is elegantly avoided in the implicit framework. Since implicit surfaces are embedded in the Euclidean space, they naturally inherit its flat metric. Thus, the ratio between the coefficients in the GAC equation remains constant and the above problems do not occur.

5 Experimental Results

We have tested our model on numerous synthetic and natural surfaces using an Intel Pentium M 1.6 GHz Notebook endowed with MATLAB®.

[1] Note, that the respective formula as stated in [9], $\partial_t \phi = -f g^{\frac{1}{2}} |\nabla_M \phi|$, is erroneous. This was acknowledged in personal communication by Prof. Ron Kimmel.

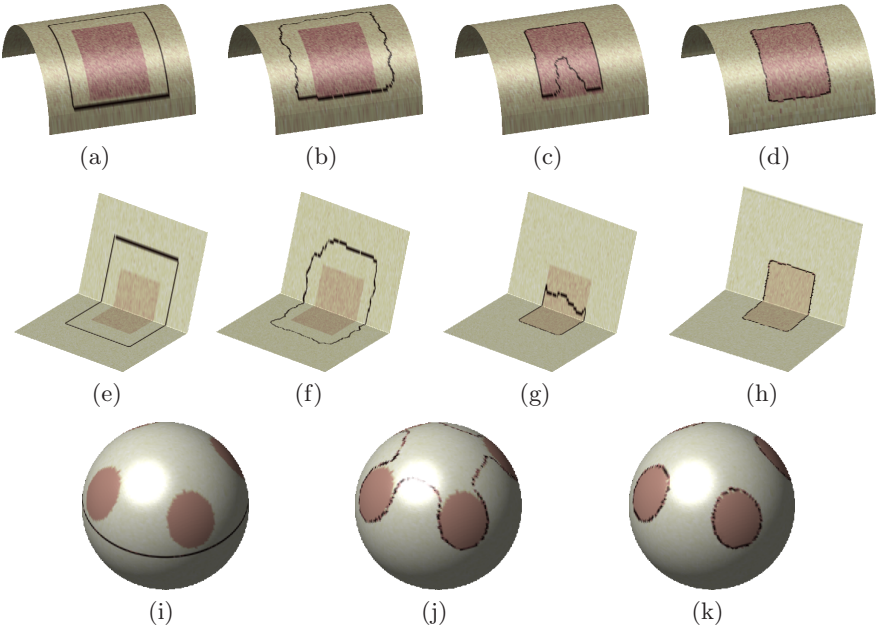

Fig. 3. (a) and (e): initial contours on a halfpipe and a notebook surface, resulting contours using Spira's parametric approach without (b,f) and with (c,g) an added balloon force. (d) and (h): results of our implicit method, (i)-(k): GAC evolution on an implicit sphere.

5.1 Comparison with Existing Methods

With our first experiments, we aimed to point out the main drawback of the GAC model [9], as analysed in Sect. 4. We could show that this drawback is overcome by our model. Both test surfaces, a halfpipe and a plane with a kink (we also call it 'notebook' surface), were parametrised as graphs $z = f(x)$. As the GAC model by Spira [9] includes no balloon force it failed in our tests with noisy texture data (see Figs. 3(b) and (f)). Then we added a balloon term to the equation (see the remarks in Sect. 4). Figures 3(c) and (g) show that the contour failed to converge to the desired object boundary, in fact it did not converge at all. In the surface areas with a larger slope – corresponding to large metric coefficients g_{ij} – the balloon term is weighted too high and thus the contour passes the object edges. We remark, that this is a systematic and reproducable effect.

Further note the following: The metric tensors of the surfaces in Fig. 3 are given by

$$(g_{ij}) = \begin{pmatrix} 1 + (\partial_x f)^2 & 0 \\ 0 & 1 \end{pmatrix} . \tag{13}$$

Table 1. Comparison of GACs using implicit and parametric surface representation: textures contain additive Gaussian noise with given variance values. The texture for the plane and the pyramid was similar to that of the notebook surface. The implicit algorithm with quadratic complexity was implemented as in [10]. BF means Balloon force, the parameterisation plane had the size of the first two dimensions of the implicit grid: e.g. $100 \times 100 \times 30$ grid means 100×100 parameterisation plane. Uniform time step for the parametric model was $\tau = 0.4$, for the implicit models we chose $\tau = 0.4$ for the plane and the pyramid and $\tau = 1.0$ otherwise. GAC parameters: $c = 0.4$ and $\mu = 1$.

			quadr. impl.		lin. impl.		parametric		
Surface	Grid size	Noise	l^∞-err.	time	l^∞-err.	time	BF	l^∞-err.	time
Plane	$100 \times 100 \times 30$	0.003	1.18	752 s	1.2	191 s	yes	1.1	171 s
Pyramid	$100 \times 100 \times 65$	none	0.91	529 s	0.89	132 s	no	0.95	257 s
Halfpipe	$130 \times 140 \times 75$	0.003	1.73	3061 s	1.77	657 s	yes/no	fails	
Notebook	150^3	0.001	1.67	2445 s	1.65	468 s	yes/no	fails	
Sphere	100^3	0.001	0.92	1474 s	0.94	360 s	yes/no	not applicable	

Thus, the metrics of the surfaces vary only in one direction and the contours pass only the edges in the respective directions while converging to the others correctly (Figs. 3(c) and (g)). The desired results, as obtained by our method with implicit surface representations, are depicted in Figs. 3(d) and (h).

These experiments and the theoretical considerations in Sect. 4 show that using the parametric framework [9] a balloon force cannot be integrated properly into the GAC model. If the metric varies, the nonuniform scaling behaviour of the equation leads systematically to changing weights in the equation and undesired segmentation results.

Let us regard Table 1: the test in the plane shows that when both the parametric and the implicit framework can be applied, the quality of the results is comparable. Then we mapped a simple texture without noise to a pyramid surface and compared Spira's method without a balloon term to ours with a balloon term. While the accuracy results were similar, our method performed faster then Spira's due to much less needed iteration steps. The experiment on a sphere illustrates a further advantage of our framework: many surfaces, including the full sphere, cannot be parametrised without singularities, thus Spira's method is not applicable (cf. his experiments in [9], where the poles of the sphere are cut). In contrast, given an implicit sphere we can evolve GAC on it easily, as shown in Figs. 3(i)-(k).

Compared to the quadratic algorithm for GAC on implicit surfaces as proposed in [11] and applied in [10], the tests in Table 1 show a considerable speedup while maintaining the accuracy. Our tests confirm that the optimized framework presented in this paper has linear complexity and processing times basically depend on the length of the evolving curve.

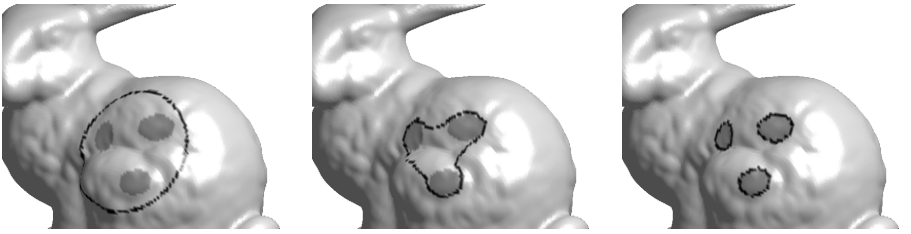

(a) (b) (c)

Fig. 4. Our GAC model applied to the Stanford bunny

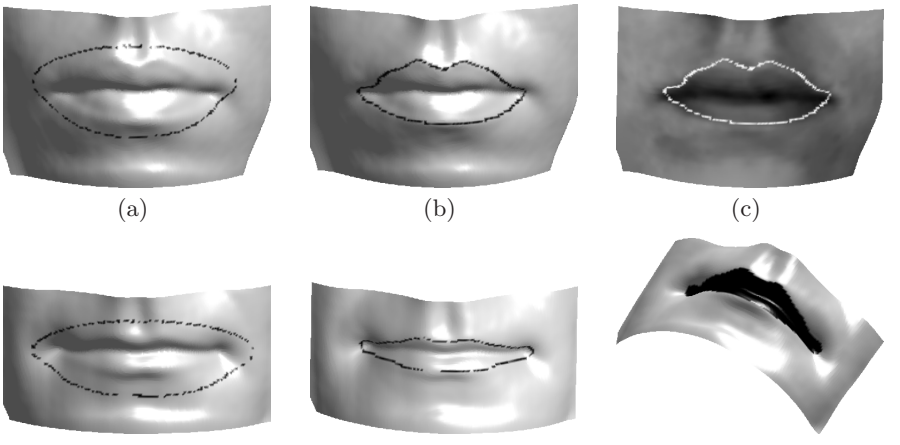

(a) (b) (c)

(d) (e) (f)

Fig. 5. 3D lip contour segmentation: initial (a,d) and final (b,e) contour, surface with added texture (c) and segmented lip area (f)

5.2 Application to 3D Face Feature Segmentation

One of the prospective application areas for our GAC model on surfaces is feature extraction from scanned 3D human faces, e.g. 3D lip contour segmentation. Feature extraction from 3D faces is important for face recognition and biometrics, yet remains a relatively young and unexplored field. However, recent research [6,7,8] emphasizes that multi-modal approaches, combining 2D and 3D information, perform better than either 2D or 3D methods alone.

As for 3D lip segmentation, we have so far tested our method with promising results on the Face Database of the MPI for Biological Cybernetics in Tuebingen, Germany [19] (see Fig. 5 for two examples and Table 2 for processing times). A detailed qualitative and quantitative evaluation is subject of our future work.

Note, that in this subsection on applications we did not compare our method to Spira et al.'s [9] as the data were noisy and a balloon force was needed. According to our analysis in Sect. 4 and our experiments in the last subsection this is not possible in the parametric framework proposed in [9].

Table 2. Performance of the proposed implicit GAC

Surface	Grid size	Parameters c	μ	Time step	Iterations	time
Bunny	$135 \times 160 \times 160$	0.4	1.0	0.4	600	480 s
Mouth1	$94 \times 143 \times 70$	0.03	0.15	1.0	900	296 s
Mouth2	$77 \times 147 \times 65$	0.03	0.15	1.0	1000	300 s
Both Eyes	$150 \times 66 \times 82$	0.035	0.25	1.0	1500	559 s
Left Eye	$150 \times 66 \times 82$	0.035	0.25	1.0	800	159 s
Right Eye	$150 \times 66 \times 82$	0.035	0.25	1.0	500	120 s
Ear	$123 \times 82 \times 114$	0.04	0.5	0.7	2000	724 s

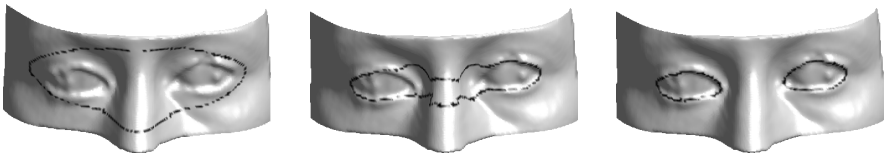

(a) (b) (c)

Fig. 6. 3D eye contour segmentation: Contour after 0, 800 and 1500 iterations

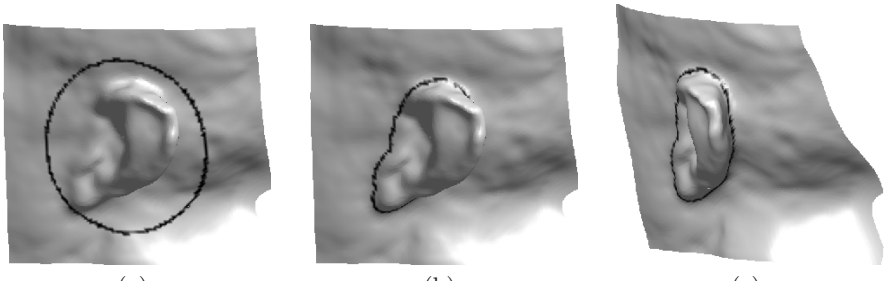

(a) (b) (c)

Fig. 7. 3D Ear segmentation: Initial (a) and final (b) contour, different view (c)

For the face experiments we defined the stopping term f in (8) by means of the mean curvature of the face surface instead of an image gradient, e.g. $f = \frac{1}{1+\tilde{H}}$ with $\tilde{H}$ a rescaled version of the mean curvature. For the lip contour experiments in particular, we designed a slightly different stopping term basing on the arccos-function. Details will be reported elsewhere. No texture data was needed to apply our method. It will be interesting to research, how our current results can be improved by incorporating the texture data.

Other face features like the eyes can be segmented by our GAC model as well. However, it is problematic that in 3D face data the eyes are often closed or of poor data quality. In Fig. 6 we show how we segmented the eyes. Note, that the results are good despite major artefacts in both eyes. The efficiency of

the algorithm can be increased by segmenting both eyes separately rather than segmenting them in one step, provided the initial contours are sufficiently close to the eyes (see Table 2).

Given a full 3D head model, 3D ear segmentation could be an interesting future application of our method (cf. [8]). Due to lack of suitable 3D head data, we applied our model to one ear of the Max Planck bust [20] (see Fig. 7). Note, that this surface is very difficult to parametrise while easier to represent implicitly using algorithms such as [21,22].

6 Conclusion

This paper has proposed a new method for evolving geodesic active contours on implicitly represented surfaces. The proposed numerical scheme reduces the computational complexity of Cheng et al.'s [11] framework from quadratic to linear. This is a necessary prerequisite for applying the model to larger datasets. We have shown theoretically and confirmed experimentally that our approach is stable on virtually every surface that can be represented implicitly. Moreover, we have identified restrictions of the published works on parametric surfaces where adding a balloon term leads to a non uniformly scaling equation causing undesired segmentation results. The strictly implicit approach and the integration of a balloon term makes our model a natural generalization of the standard GAC to curved surfaces. Initial results on 3D human face data sets indicate that our 3D active contour model is a promising avenue for 3D face feature segmentation.

Acknowledgements

M. Krueger's research is funded by The New Zealand International Doctoral Research Scholarship and The University of Auckland International Doctoral Scholarship.

References

1. Bertalmio, M., Memoli, F., Cheng, L.T., Sapiro, G., Osher, S.: Variational problems and partial differential equations on implicit surfaces: Bye bye triangulated surfaces? In: Osher, S., Paragios, N. (eds.) Geometric Level Set Methods in Imaging, Vision, and Graphics, pp. 381–398. Springer, New York (2003)
2. Bogdanova, I., Bresson, X., Thiran, J., Vandergheynst, P.: Scale-space analysis and active contours for omnidirectional images. IEEE Trans. Image Processing 16, 1888–1901 (2007)
3. Kimmel, R.: Intrinsic scale space for images on surfaces: The geodesic curvature flow. Graphical Models And Image Processing 59, 365–372 (1997)
4. Wu, C., Deng, J., Zhu, W., Chen, F.: Inpainting images on implicit surfaces. In: Proc. Pacific Graphics Conf. 2005, University of Macau, pp. 142–144 (2005)
5. do Carmo, M.: Differential Geometry of Curves and Surfaces. Prentice-Hall, Englewood Cliffs (1976)

6. Chang, K.I., Bowyer, K.W., Flynn, P.J.: An evaluation of multimodal 2D+3D face biometrics. IEEE Trans. Pattern Anal. Machine Intell. 27, 619–624 (2005)
7. Tsalakanidou, F., Malassiotis, S., Strintzis, M.G.: Integration of 2D and 3D images for enhanced face authentication. In: Proc. IEEE 6th Int. Conf. Automatic Face and Gesture Recognition (FG 2004), Seoul, Korea, pp. 266–271 (2004)
8. Yan, P., Bowyer, K.W.: Biometric recognition using 3D ear shape. IEEE Trans. Pattern Anal. Machine Intell. 29, 1297–1308 (2007)
9. Spira, A., Kimmel, R.: Geometric curve flows on parametric manifolds. J. Comp. Physics 223, 235–249 (2007)
10. Krueger, M., Delmas, P., Gimel'farb, G.: Towards feature extraction on implicit surfaces using geodesic active contours. In: Proc. Image and Vision Computing New Zealand, Hamilton, New Zealand, pp. 294–299 (2007)
11. Cheng, L.T., Burchard, P., Merriman, B., Osher, S.: Motion of curves constrained on surfaces using a level-set approach. J. Comp. Physics 175, 604–644 (2002)
12. Kass, M., Witkin, A., Terzopoulos, D.: Snakes: Active contour models. Int. J. Computer Vision 1, 321–331 (1987)
13. Caselles, V., Kimmel, R., Sapiro, G.: Geodesic active contours. Int. J. Computer Vision 22, 61–79 (1997)
14. Osher, S., Sethian, J.: Fronts propagating with curvature-dependent speed: algorithms based on Hamilton-Jacobi formulations. J. Comp. Physics 79, 12–49 (1988)
15. Jin, H., Yezzi, A., Soatto, S.: Mumford-Shah on the move: Region-based segmentation on deforming manifolds with application to 3D reconstruction of shape and appearance from multi-view images. J. Math. Imaging Vision 29, 219–234 (2007)
16. Osher, S., Fedkiw, R.: Level sets methods and dynamic implicit surfaces. Springer, New York (2003)
17. Peng, D., Merriman, B., Zhao, H.K., Osher, S., Kang, M.: A PDE based fast local level set method. J. Comp. Physics 155, 410–438 (1999)
18. Sethian, J.: A fast marching level set method for monotonically advancing fronts. Proc. Nat. Acad. Sciences 93, 1591–1595 (1996)
19. Troje, N., Buelthoff, H.H.: Face recognition under varying poses: The role of texture and shape. Vision Research 36, 1761–1771 (1996)
20. Princeton suggestive contours gallery, `http://www.cs.princeton.edu/gfx/prog/sugcon/models/`
21. Carr, J.C., Beatson, R.K., Cherrie, J.B., et al.: Reconstruction and representation of 3D objects with radial basis functions. In: Proc. SIGGRAPH 2001, Los Angeles, USA, pp. 67–76 (2001)
22. Mauch, S.: A fast algorithm for computing the closest point and distance transform. Technical Report, California Institute of Technology (2000)