

Integrating Shape and Texture in Deformable Models: from Hybrid Methods to Metamorphs

Dimitris Metaxas¹, Xiaolei Huang¹, and Ting Chen²

¹ Division of Computer and Information Science,
Rutgers, the State University of New Jersey, New Brunswick, NJ, USA

² Department of Radiology
NYU Medical School, New York, NY, USA

Abstract. In traditional shape-based deformable models, the external image forces come primarily from edge or gradient information. Such reliance on edge information, however, makes the models prone to get stuck in local minima due to image noise and various other image artifacts. Integrating region statistics constraints has been a centerpiece of the efforts toward more robust, well-behaved deformable models in boundary extraction and segmentation. In this chapter, we review previous work on the loose coupling of boundary and region information in two major classes of deformable models: the parametric models and the geometric models. Then, we propose a new class of deformable shape and texture models, which we term “Metamorphs”. The novel formulation of the Metamorph models tightly couples shape and interior texture and the dynamics of the models are derived in a unified manner from both boundary and region information in a variational framework.

Keywords

Metamorphs, deformable models, implicit representation, shape, texture, calculus of variation, distance transform, nonparametric region statistics, hybrid segmentation methods

1 Introduction

Automated image segmentation is a fundamental problem in computer vision and medical image analysis applications. Object texture, image noise, intensity inhomogeneity and variations in lighting, to name a few, add to the problem complexity. To address these difficulties, deformable model-based segmentation methods have been extensively studied and widely used, with promising results.

Deformable models are curves or surfaces that move under the influence of internal smoothness and external image forces. In the literature, there are two major classes of deformable models. The first is the parametric (explicit) deformable models that explicitly represent deformable curves and surfaces in their parametric form during the segmentation process. Examples are Active Contour Models [11] and their extensions in both 2D and 3D [14, 22, 6, 13, 7,

18, 25]. The evolution of these parametric models is derived either in a energy-minimization process [11, 26] or through a dynamic-force formulation [6]. The energy-minimization formulation has the advantage that its solution satisfies a minimum principle; while the dynamic force formulation provides the flexibility of applying different types of external forces onto the deformable model. The external forces can be potential forces such as image forces, non-potential forces such as balloon forces, and the combination of both. The other class of deformable models is the geometric (implicit) deformable models [3, 12, 29, 28, 4]. These models represent curves and surfaces implicitly as the level set of a higher-dimensional scalar function [21, 16], and the model evolution is based on the theory of curve evolution, with speed function specifically designed to incorporate image information. Comparing the two classes of deformable models, the parametric deformable models have a compact representation, and allow fast implementation, while the geometric deformable models can handle topological changes naturally.

Although the parametric and geometric deformable models differ both in their formulations and in their implementations, both classes use primarily edge (image gradient) information to derive external image forces to drive a shape-based model. Such reliance on edge information, however, makes the models sensitive to image noise and various other image artifacts. For instance, a model may leak through small or large gaps on the object boundary, or it may get stuck in local minima due to spurious edges inside the object or clutter around the true boundary.

To address these limitations, there have been significant efforts in the literature to integrate region information into both parametric [19, 30] and geometric deformable models [17, 24, 23]. The integration frameworks however, are still imperfect. In the case of parametric models, region information and boundary information are often treated separately in different energy minimization processes, thus parameters of region intensity statistics can not be updated simultaneously with the boundary shape parameters. In the case of geometric models, the integrations are mostly based on solving reduced cases of the minimal partition problem in the Mumford and Shah model for segmentation [15]. Variational frameworks are proposed to unifying boundary and region-based information sources, and level set approaches are used to implement the resulting PDE systems. However, these frameworks assume piecewise constant, or Gaussian intensity distributions within each partitioned region. This limits their applicability and robustness in finding objects whose interiors have high noise level, intensity inhomogeneity, and/or complex multi-modal intensity distributions.

In this chapter, we focus on presenting the work from our group on the integration of region statistics constraints into shape-based deformable models, which includes: (1) a hybrid framework that loosely couples a region-based module and a boundary deformable model-based module, and (2) *Metamorphs*, a recently developed new class of deformable models that possess both shape and interior texture and integrate boundary and region information in a unified manner within a variational framework.

In [5], we proposed a hybrid segmentation framework which integrates a region-based segmentation module driven by Gibbs prior models, a boundary-based module using deformable models and the marching cubes method which connects these two modules. The region-based and boundary based modules work recursively: The region segmentation results are used to initialize the deformable model and the deformable fitting results are used to update the parameters of the region segmentation. This way, the two modules can help each other out of local minima. The quality of the segmentation output also improves when we update the Gibbs model’s parameters using more accurate region and boundary information at the end of each iteration. To accommodate 3D segmentation applications, we integrate the marching cubes method into our method, which can construct deformable meshes based on 3D binary masks.

One limitation in the hybrid framework, however, is that the region information and the boundary/shape information are still treated separately instead of being integrated in driving model deformations. To utilize information from both sources in a unified manner, we have developed, recently, a new class of deformable models called “Metamorphs” [9].

Metamorphs integrate dynamically shape and interior texture. The resulting lagrangian formulation is derived from both boundary and region information based on a novel variational framework. These new models bridge the gap between parametric and geometric deformable models by borrowing the best features of both worlds. The model shapes are embedded in a higher dimensional space of distance transforms, thus represented by distance map “images”. (This is similar to the implicit shape representation in geometric level-set based models). The model deformations are efficiently parameterized using a space warping technique, the cubic B-spline based Free Form Deformations (FFD) [1, 2, 10].³ The interior intensity statistics of the models are captured using nonparametric kernel-based approximations, which can represent complex multi-modal distributions. When finding object boundaries in images, the dynamics of the Metamorph models are derived from an energy functional consisting of both edge (which encodes gradient information) and region intensity energy terms. In our formulation, both types of energy terms are differentiable with respect to the model deformation parameters. This allows for a unified gradient-descent based deformation parameter updating paradigm using both boundary and region information. Furthermore, our Metamorph model deformations are constrained in such way that the interior statistics of the model after deformation is consistent with the statistics learned from the past history of the model interiors. A Metamorph model can be initialized far-away from the object boundary and efficiently converge to an optimal solution. The proposed energy functional enables the model to pass small spurious edges and prevents it from leaking through large boundary gaps, hence makes the boundary finding robust to image noise and inhomogeneity.

³ Note that we separate the shape representation, which is implicit in a higher dimension, and model deformation, which is explicitly parameterized by FFD.

In the remainder of the chapter, we will first present our hybrid segmentation framework and then the new form of deformable shape and texture models - Metamorphs.

2 Hybrid Segmentation Method

In the framework we proposed in [5], we segment an object as follows. First we use the Gibbs model to get a rough binary mask of the object. Then we use the marching cubes method to construct the deformable mesh and make the mesh deform to fit the object surface using the gradient information. The Gibbs parameters need to be updated from iteration to iteration to improve the segmentation results. By doing so, we integrated the region information into deformable models. In the following, we present the modules that comprise the hybrid segmentation approach.

2.1 Gibbs Models

Most medical images are Markov Random Field images, that is, the statistics of a pixel in the medical image are related to the statistics of pixels in its neighborhood. According to the Equivalence Theorem proved by Hammersley and Clifford [8], a Markov Random Field is equivalent to a Gibbs field under certain restrictions. Therefore the joint distribution of a medical image with MRF property can be written in the Gibbsian form as follows.

$$\Pi(X) = Z^{-1} \exp(-H(X)) \quad (1)$$

where \mathbf{X} is the set of all possible configurations of the image X , z is an image in the set of \mathbf{X} , $Z = \sum_{z \in \mathbf{X}} \exp(-H(z))$ is a normalizing factor, and $H(X)$ is the energy function of image X . The local and global properties of MRF images are incorporated into the model by designing an appropriate energy function $H(X)$ and minimizing it. The lower the value of the energy function, the better the image fits to the prior distribution. Therefore the segmentation procedure corresponds to the minimization of the energy function.

$$H_{prior}(X) = H_1(X) + H_2(X) \quad (2)$$

where $H_1(X)$ models the piecewise pixel homogeneity statistics and $H_2(X)$ models the object boundary continuity. In general, the homogeneity term $H_1(X)$ has a smoothing effect on pixels inside the object and will leave boundary features beyond the threshold unchanged. The boundary continuity term in the energy function $H_2(X)$ has the following form:

$$H_2(X) = \vartheta_2 \sum_{s \in X} \sum_{i=1}^N W_i(s) \quad (3)$$

where s is a pixel, ϑ_2 is the weight term for the boundary continuity, N is the number of local configurations, and $W_i(s)$ are weight functions (also called the




Fig. 1. Clique definitions: cliques can be classified into clique types of a) smooth boundary, b) smooth boundary with angle, c) smooth boundary in diagonal direction, d) object interior, e) outside the object, f) irregular boundaries or noisy regions. Pixels labelled 1 are in the object, while pixels labelled 0 are out of the object.

potential functions) of local configurations. In our model, the potential functions are defined on a neighborhood system based on cliques with the size of 3 by 3 pixels. There are altogether 2^9 possible local configurations in a clique including 3 by 3 pixels. We can classify them into 6 clique types. Among these 6 types, three of them contain configurations at smooth boundaries (Fig.1.a, .b, .c), one type for the homogeneous region inside (Fig.1.d) and one for such region outside (Fig.1.e) the object respectively, and one clique type includes all local configurations that lead to noisy regions or irregular boundaries (Fig.1.f). Cliques that belong to the same clique type share the same potential value. We assign lower potential values to clique configurations that are located at smooth and continuous boundaries. Therefore, when we minimize $H_2(X)$, pixels in the image (especially those near the boundary) will alter their intensities to form clique configurations of lower potentials. These alternations make the currently estimated boundaries smoother, the weak boundaries stronger, and extend boundaries into image regions without strong gradient information. Hence the minimization of the energy function will lead to continuous and smooth object boundaries.

We use the Bayesian framework to get a MAP estimation of the object region. In a Bayesian framework, the segmentation problem can be formulated as the maximization of the posterior probability $P(X|Y)$, which can also be written as an energy functional:

$$H_{posterior}(X, Y) = H_{prior}(X) + H_{observation}(X, Y) \quad (4)$$

where $H_{observation}(X, Y)$ is the constraint from the observation of the original image. Using $H_{posterior}(X, Y)$ instead of $H_{prior}(X)$ in the energy minimization, we get a MAP estimation of the object region. The constraint of the observation will compete with the prior distribution during the minimization process. Hence the result of the minimization process will still be close enough to the original observation, while important image features, such as irregular edges, will be kept regardless of the prior distribution.

The output of the Gibbs prior model includes region information so that when its output is used to initialize the geometric form of the deformable, the region information will be passed to the deformable.

2.2 Deformable models in the Hybrid Framework

We use Gibbs models and the marching cubes method to construct the geometry of the deformable model, i.e, a deformable surface close to the object surface. Then we write the deformable model dynamics in the form of the first order Lagrangian equation:

$$\dot{\mathbf{d}} + \mathbf{K}\mathbf{d} = \mathbf{f}_{ext} \quad (5)$$

where $\dot{\mathbf{d}} = \frac{\partial \mathbf{X}}{\partial t}$. \mathbf{K} is the stiffness matrix. \mathbf{f}_{ext} is the external force.

According to equation (5), the deformable model deforms under the effect of the internal force $\mathbf{K}\mathbf{d}$ and the external force. The internal force keeps the deformable model surface smooth and continuous during its deformation. If the object boundary in the image to be segmented is weak, the internal force will act as a surface constraint that prevents the model from being trapped into local minima or overflowing beyond the boundary. The external force will lead the model to the object surface using image information such as the gradient.

In our framework, we use the second order derivative gradient as the external force. It is defined as:

$$f_G(x, y, z) = -\nabla P(x, y, z) = -\nabla(w_e |\nabla[G_\sigma(x, y, z) * I(x, y, z)]|)^2 \quad (6)$$

where $I(x, y, z)$ is the original image, w_e is a positive weighting parameter, $G_\sigma(x, y, z)$ is a three dimensional Gaussian function with standard deviation σ , ∇ is the gradient operator, and $*$ is the convolution operator. We use the Gaussian filter to blur the original image in order to remove small noisy regions and expand the effective range of the gradient-derived force. In a second order gradient flow field, all gradient vectors point to the location of edge features so that they can lead the model to the object surface directly. During the fitting process, we calculate the dot product of the second order gradient vector and the normal vector at every node on the deformable surface. It yields a positive value if the model node locates inside the edge feature and a negative value if the model node locates outside. We can define the magnitude of the external force as the magnitude of the dot product and the direction of the force vector as the direction of the normal vector at the node.

We now can calculate the derivative of displacements of every node on the deformable model surface using Eqn. (6). The displacements will then be updated using the Euler equation:

$$\mathbf{d}_{new} = \dot{\mathbf{d}} \cdot \Delta t + \mathbf{d}_{old} \quad (7)$$

where Δt is the time step. The deformation stops when the forces equilibrate or vanish.

2.3 Integration of Deformable Models and Gibbs Models

Fig. 2 shows internal modules and the data flow of our 3D hybrid segmentation framework. In the first iteration of the recursive hybrid framework, the parameters of the Gibbs prior models are set to default values. Using the segmentation

result of the deformable model in the current iteration, we update the Gibbs prior parameters before restarting the Gibbs models in the following iterations to improve their segmentation performance. Besides updating regional parameters such as the mean intensity and the standard deviation of the object, we also update potentials of local configurations in Eqn. (3). The clique potentials of the Gibbs Prior model are set to be proportional to the number of appearances of each type of cliques in the deformable model segmented binary image.




Fig. 2. Flow-Chart for 3D-segmentation hybrid framework.




Fig. 3. Segmentation of a tumor in the brain from MR image, (a) the original image; (b & c) the Gibbs model segmentation result in the first and second iterations; (d) the final segmentation result of the hybrid framework; (e) the segmentation result overlaid upon the original image; (f) the initial deformable surface; (g, h) 2 views of the final segmentation result in 3D. Data courtesy of Prof. Kikinis’s group at Harvard University.

We illustrate our hybrid segmentation framework by applying it to segment the tumor region in a 3D MR image volume of human brain (See Fig. 3). Fig. 3(a) shows one slice of the volume. The image volume size is 256 by 256 by 32

pixels (preprocessing has been applied to remove slices that do not contain the structure of interest). We use 32 2D Gibbs Prior models to create a 3D binary mask for the tumor region. The initial edge threshold is set to 6, the potential weight for smooth boundaries are set to 0.0, and the potential for other local configurations are set to 5.0. We then use the marching cube method to create a surface mesh for the deformable model to begin with. During the deformable model fitting, the time step is set to 0.07, and the gradient magnitude w_e is 1.0. The hybrid segmentation process stops after two iterations. Fig. 3(d) shows the final segmentation result of the hybrid framework. For quality evaluation purposes, we overlay the segmentation result onto the original image 3(a) as in Fig. 3(e). We show the initial deformable mesh surface constructed by the marching cube method in Fig. 3(f), and the 3D reconstruction of the tumor region based on final segmentation result in Fig. 3(g), (h). Notice that the segmentation result of the Gibbs model is improved by using updated parameters. The fact that in Fig. 3(g) and (h) the deformable model fits well at concavities and convexities proves that our hybrid framework has a good performance in segmenting complex object surfaces. The total segmentation time is about 6 minutes for 2 iterations, which is much shorter than the method described in [27].

3 Metamorphs: Deformable Shape and Texture Models

A limitation in the hybrid segmentation framework introduced in section 2 is that, the region-based module and the boundary-based module are used separately, thus the information from both sources are not integrated during the evolution of a deformable model. Furthermore, the region-based module produces an initialization mesh to start a deformable model, which makes the final segmentation result highly dependent on this initialization. To address these limitations, we present our recent work [9] on a new class of deformable shape and texture models, which we call “Metamorphs”. The formulation of Metamorphs naturally integrates both shape and interior texture, and the model dynamics are derived coherently from both boundary and region information during the whole course of model evolution in a common variational framework.

3.1 The Metamorphs Model representations

The Model’s Shape Representation The model’s shape is embedded implicitly in a higher dimensional space of distance transforms. The Euclidean distance transform is used to embed an evolving model as the zero level set of a higher dimensional distance function. In order to facilitate notation, we consider the 2D case. Let $\Phi : \Omega \rightarrow R^+$ be a Lipschitz function that refers to the distance transform for the model shape \mathcal{M} . The shape defines a partition of the domain: the region that is enclosed by \mathcal{M} , $[\mathcal{R}_{\mathcal{M}}]$, the background $[\Omega - \mathcal{R}_{\mathcal{M}}]$, and on the model, $[\partial\mathcal{R}_{\mathcal{M}}]$ (In practice, we consider a narrow band around the model \mathcal{M} in the image domain as $\partial\mathcal{R}_{\mathcal{M}}$). Given these definitions the following implicit shape

representation is considered:

$$\Phi_{\mathcal{M}}(\mathbf{x}) = \begin{cases} 0, & \mathbf{x} \in \partial\mathcal{R}_{\mathcal{M}} \\ +ED(\mathbf{x}, \mathcal{M}) > 0, & \mathbf{x} \in \mathcal{R}_{\mathcal{M}} \\ -ED(\mathbf{x}, \mathcal{M}) < 0, & \mathbf{x} \in [\Omega - \mathcal{R}_{\mathcal{M}}] \end{cases}$$

where $ED(\mathbf{x}, \mathcal{M})$ refers to the min Euclidean distance between the image pixel location $\mathbf{x} = (x, y)$ and the model \mathcal{M} .

Such treatment makes the model shape representation a distance map “image”, which greatly facilitates the integration of boundary and region information. This shape representation in 3D is similarly defined in a volumetric embedding space.

The Model’s Deformations The deformations that Metamorph models can undergo are defined using a space warping technique, the Free Form Deformations (FFD) [20]. The essence of FFD is to deform an object by manipulating a regular control lattice F overlaid on its volumetric embedding space. In Metamorphs, we consider an Incremental Free Form Deformations (IFFD) formulation using the cubic B-spline basis [10].

Let us consider a regular lattice of control points

$$F_{m,n} = (F_{m,n}^x, F_{m,n}^y); \quad m = 1, \dots, M, \quad n = 1, \dots, N$$

overlaid to a region $\Gamma_c = \{\mathbf{x}\} = \{(x, y) | 1 \leq x \leq X, 1 \leq y \leq Y\}$ in the embedding space that encloses the model in its object-centered coordinate system. Let us denote the initial configuration of the control lattice as F^0 , and the deforming control lattice as $F = F^0 + \delta F$. Under these assumptions, the incremental FFD parameters, which are also the deformation parameters for the model, are the deformations of the control points in both directions (x, y) :

$$\mathbf{q} = \{(\delta F_{m,n}^x, \delta F_{m,n}^y)\}; \quad (m, n) \in [1, M] \times [1, N]$$

The deformed position of a pixel $\mathbf{x} = (x, y)$ given the deformation of the control lattice from F^0 to F , is defined in terms of a tensor product of Cubic B-spline polynomials:

$$D(\mathbf{q}; \mathbf{x}) = \mathbf{x} + \delta D(\mathbf{q}; \mathbf{x}) = \sum_{k=0}^3 \sum_{l=0}^3 B_k(u) B_l(v) F_{i+k, j+l}^0 + \delta F_{i+k, j+l} \quad (8)$$

where $i = \lfloor \frac{x}{X} \cdot (M-1) \rfloor + 1$, $j = \lfloor \frac{y}{Y} \cdot (N-1) \rfloor + 1$. The terms of the deformation component refer to:

- $\delta F_{i+l, j+l}$, $(k, l) \in [0, 3] \times [0, 3]$ are the deformations of pixel \mathbf{x} ’s (sixteen) adjacent control points,
- $B_k(u)$ is the k^{th} basis function of a Cubic B-spline, defined by:

$$B_0(u) = (1-u)^3/6, \quad B_1(u) = (3u^3 - 6u^2 + 4)/6 \\ B_2(u) = (-3u^3 + 3u^2 + 3u + 1)/6, \quad B_3(u) = u^3/6$$

with $u = \frac{x}{X} \cdot (M-1) - \lfloor \frac{x}{X} \cdot (M-1) \rfloor$. $B_l(v)$ is similarly defined.




Fig. 4. The Left Ventricle Endocardium segmentation. (1) Initial model. (2) Intermediate result. (3) Final converged result. (a) The evolving model drawn in colored lines (blue or red) on original image. (b) Interior of the evolving model. (c) The intensity p.d.f of the model interior. The X axis is the intensity value in the range of $[0, 255]$ and the Y axis is the probability value in the range of $[0, 1]$. (d) The image probability map based on the p.d.f of the model interior.

– $\delta D(\mathbf{q}; \mathbf{x}) = \sum_{k=0}^3 \sum_{l=0}^3 B_k(u)B_l(v)\delta F_{i+k,j+l}$ is the incremental deformation for pixel \mathbf{x} .

The extension of the models to account for deformations in 3D is straightforward, by using control lattices in the 3D space and a 3D tensor product of B-spline polynomials.

The Model’s Texture Rather than using traditional statistical parameters (such as mean and variance) to approximate the intensity distribution of the model interior, we model the distribution using a nonparametric kernel-based method. The nonparametric approximation is differentiable, more generic and can represent complex multi-modal intensity distributions.

Suppose the model is placed on an image I , the image region bounded by current model $\Phi_{\mathcal{M}}$ is $\mathcal{R}_{\mathcal{M}}$, then the probability of a pixel’s intensity value i being consistent with the model interior intensity can be derived using a Gaussian kernel as:

$$\mathbf{P}(i|\Phi_{\mathcal{M}}) = \frac{1}{V(\mathcal{R}_{\mathcal{M}})} \iint_{\mathcal{R}_{\mathcal{M}}} \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(i-I(\mathbf{y}))^2}{2\sigma^2}} d\mathbf{y} \quad (9)$$

where $V(\mathcal{R}_{\mathcal{M}})$ denotes the volume of $\mathcal{R}_{\mathcal{M}}$, and σ is a constant specifying the width of the gaussian kernel.

Using this nonparametric approximation, the intensity distribution of the model interior gets updated automatically while the model deforms. The initialization of the model texture is flexible. We can either start with a small model inside the texture region to be segmented, or use supervised learning to specify the desired texture a Priori. One example of the model interior texture representation can be seen in [Fig. (4)]. In the figure, we show the zero level set of the current model $\Phi_{\mathcal{M}}$ in colored lines [Fig. (4).a], the model interior region $\mathcal{R}_{\mathcal{M}}$ [Fig. (4).b], the probability density function (p.d.f.) for the intensity of current model interior $\mathbf{P}(i|\Phi_{\mathcal{M}})$ for $i = 0, \dots, 255$ [Fig. (4).c], and the probability map of every pixel’s intensity in the image according to the model interior distribution [Fig. (4).d].

3.2 The Metamorph Dynamics

The motion of the model is driven by both boundary (edge) and region (intensity) energy terms derived from the image. The overall energy functional E consists of two parts – the shape data terms E_S , and the intensity data terms E_I :

$$E = E_S + kE_I \quad (10)$$

where k is a constant balancing the contribution of the two parts. Next, we derive the shape and intensity data terms respectively.

The Shape Data Terms We encode the gradient information of an image using a “shape image” Φ , which is derived from the un-signed distance transform of the edge map of the image. In [Fig. (5).c], we can see the “shape image” of an example MR heart image.

To evolve a Metamorph model toward image edges, we define two shape data terms – an interior term E_{S_i} and a boundary term E_{S_b} :

$$E_S = E_{S_i} + aE_{S_b} \quad (11)$$

In the interior shape data term of the model, we aim to minimize the Sum-of-Squared-Differences between the implicit shape representation values in the model interior and the underlying “shape image” values at corresponding deformed positions. This can be written as:

$$E_{S_i} = \frac{1}{V(\mathcal{R}_{\mathcal{M}})} \iint_{\mathcal{R}_{\mathcal{M}}} (\Phi_{\mathcal{M}}(\mathbf{x}) - \Phi(D(\mathbf{q}; \mathbf{x})))^2 d\mathbf{x} \quad (12)$$

During optimization, this term will deform the model along the gradient direction of the underlying “shape image”. Thus it will expand or shrink the model accordingly, serving as a two-way balloon force without explicitly introducing such forces, and making the attraction range of the model large.

To make the model deformation more robust to small spurious edges detected within an object due to texture, we consider a separated boundary shape data




Fig. 5. The effect of small spurious edges inside the object of interest (endocardium of the Left Ventricle) on the “shape image”. (a) The original MR image. (b) The edge map of the image. (c) The derived “shape image”, with edges points drawn in yellow. Note the effect of the small spurious edges on the “shape image” inside the object.




Fig. 6. The boundary shape data term constraints at small gaps in the edge map. (a) Original Image. (b) The edge map, note the small gap inside the red square region. (c) The “shape image”. (d) Zoom-in view of the region inside the red square. The numbers are the “shape image” values at each pixel location. The red dots are edge points, the blue squares indicate a path favored by the boundary term for a Metamorph model.

term, which allows higher weights for pixels in a narrow band around the model boundary $\partial\mathcal{R}_M$.

$$E_{S_b} = \frac{1}{V(\partial\mathcal{R}_M)} \iint_{\partial\mathcal{R}_M} (\Phi(D(\mathbf{q}; \mathbf{x})))^2 d\mathbf{x} \quad (13)$$

Intuitively, this term will encourage the deformation that maps the model boundary to the image edge locations where the underlying “shape image” distance values are as small (or as close to zero) as possible. One additional advantage of this term is that, at an edge with small gaps, this term will constrain the model to go along the “geodesic” path, which coincides with the smooth shortest path connecting the two open ends of a gap. This behavior can be seen from [Fig. (6)]. Note that at a small gap of the edge map, the boundary term will favor a path with the smallest accumulative distance values to the edge points.




Fig. 7. Deriving the “region of interest” intensity data term. (a) The model shown (in yellow) on the original image. (b) The intensity probability map based on the model interior statistics. (c) The region of interest (ROI) derived from the thresholded probability map. The threshold is the mean probability over the entire image. (d) The “shape image” encoding boundary information of the ROI.

The Intensity Data Terms In our current framework, the intensity energy function E_I consists of two intensity data terms – a “Region Of Interest” (ROI) term E_{I_r} , and a Maximum Likelihood term E_{I_m} :

$$E_I = E_{I_r} + bE_{I_m} \quad (14)$$

In the “Region Of Interest” (ROI) term E_{I_r} , we aim to evolve the model toward the boundary of current region of interest, which is determined based on current model interior intensity distribution. Given a model \mathcal{M} on image I [Fig. (7).a], we first compute the image intensity probability map P_I [Fig. (7).b], based on the model interior intensity statistics (see section 3.1). Then a small threshold (typically the mean probability over the entire image domain) is applied on P_I to produce a binary image BP_I , in which pixels with probabilities higher than the threshold have value 1. Morphological operations are used to fill in small holes in BP_I . We then take the connected component on this binary image overlapping the model as current region of interest (ROI). Suppose the binary mask of this ROI is BI_r [Fig. (7).c], we encode its boundary information by computing the “shape image” of BI_r , which is the un-signed distance transform of the region boundary [Fig. (7).d]. Denote this “shape image” as Φ_r , the ROI intensity data term is defined as follows:

$$E_{I_r} = \frac{1}{V(\mathcal{R}_{\mathcal{M}})} \iint_{\mathcal{R}_{\mathcal{M}}} (\Phi_{\mathcal{M}}(\mathbf{x}) - \Phi_r(D(\mathbf{q}; \mathbf{x})))^2 d\mathbf{x} \quad (15)$$

This ROI intensity data term is the most effective in countering the effect of small spurious edges inside the object of interest (e.g. in Figs. (5,9)). It also provides implicit balloon forces to quickly deform the model toward object boundary.

To achieve better convergence when the model gets close to the object boundary, we design another Maximum Likelihood (ML) intensity data term that constrains the model to deform toward areas where the pixel probabilities of belonging to the model interior intensity distribution are high. This ML term is formalized by maximizing the log-likelihood of pixel intensities in a narrow band




Fig. 8. Segmentation of the Endocardium of the Left Ventricle in a MR image with a large portion of the object boundary edge missing. (1.a) The original image. (1.b) The “shape image” derived from edge map. (1.c) The intensity probability map based on the initial model. (2.a) Initial model (zero level set shown in blue). (2.b) Intermediate model (zero level set shown in red). (2.c) converged model.

around the model after deformation:

$$\begin{aligned}
 E_{I_m} &= -\frac{1}{V(\partial\mathcal{R}_{\mathcal{M}})} \iint_{\partial\mathcal{R}_{\mathcal{M}}} \log \mathbf{P}(I(D(\mathbf{q}; \mathbf{x})) | \Phi_{\mathcal{M}}) d\mathbf{x} \\
 &= -\frac{1}{V(\partial\mathcal{R}_{\mathcal{M}})} \iint_{\partial\mathcal{R}_{\mathcal{M}}} \left[\log \frac{1}{V(\mathcal{R}_{\mathcal{M}})} + \log \frac{1}{\sqrt{2\pi}\sigma} \right. \\
 &\quad \left. + \log \iint_{\mathcal{R}_{\mathcal{M}}} e^{-\frac{(I(D(\mathbf{q}; \mathbf{x})) - I(\mathbf{y}))^2}{2\sigma^2}} d\mathbf{y} \right] d\mathbf{x} \tag{16}
 \end{aligned}$$

During model evolution, when the model is still far away from object boundary, this ML term generates very little forces to influence the model deformation. When the model gets close to object boundary, however, the ML term generates significant forces to prevent the model from leaking through large gaps (e.g. in Fig. 8), and help the model to converge to the true object boundary.

3.3 Model Evolution

In our formulations above, both shape data terms and intensity data terms are differentiable with respect to the model deformation parameters \mathbf{q} , thus a unified gradient-descent based parameter updating scheme can be derived using both boundary and region information. Based on the definitions of the energy functions, one can derive the following evolution equation for each element \mathbf{q}_i in the model deformation parameters \mathbf{q} :

$$\frac{\partial E}{\partial \mathbf{q}_i} = \left(\frac{\partial E_{S_i}}{\partial \mathbf{q}_i} + a \frac{\partial E_{S_b}}{\partial \mathbf{q}_i} \right) + k \left(\frac{\partial E_{I_r}}{\partial \mathbf{q}_i} + b \frac{\partial E_{I_m}}{\partial \mathbf{q}_i} \right) \tag{17}$$

The detailed derivations for each term can be found in [9].




Fig. 9. The tagged MR heart image. (1.a) The original image. (1.b) The edge map. (1.c) The edge points overlaid on original image. (1.d) The “shape image”. (2) Initial model. (3) Intermediate result. (4) Final model (after 50 iterations). (2-4)(a) The evolving model. (2-4)(b) The model interior. (2-4)(c) The model interior intensity probability density. (2-4)(d) The intensity probability map of the image based on the p.d.f in (c).

3.4 The Model Fitting Algorithm and Experimental Results

The overall model fitting algorithm consists of the following steps:

1. Initialize the deformation parameters \mathbf{q} to be \mathbf{q}^0 , which indicates no deformation.
2. Compute $\frac{\partial E}{\partial \mathbf{q}_i}$ for each element \mathbf{q}_i in the deformation parameters \mathbf{q} .
3. Update the parameters $\mathbf{q}'_i = \mathbf{q}_i - \lambda \cdot \frac{\partial E}{\partial \mathbf{q}_i}$.
4. Using the new parameters, compute the new model $\mathcal{M}' = D(\mathbf{q}'; \mathcal{M})$.
5. Update the model. Let $\mathcal{M} = \mathcal{M}'$, re-compute the implicit representation of the model $\Phi_{\mathcal{M}}$, and the new partitions of the image domain by the new model: $[\mathcal{R}_{\mathcal{M}}]$, $[\Omega - \mathcal{R}_{\mathcal{M}}]$, and $[\partial \mathcal{R}_{\mathcal{M}}]$. Also re-initialize a regular FFD control lattice to cover the new model, and update the “region of interest” shape image ϕ_r based on the new model interior.
6. Repeat steps 1-5 until convergence.

In the algorithm, after each iteration, both shape and interior intensity statistics of the model get updated based on the model dynamics, and deformation parameters get re-initialized for the new model. This allows continuous, both large-scale and small-scale deformations for the model to converge to the energy minimum.

Some examples of using our Metamorph models for boundary finding in images have been shown in [Fig. (4)] and [Fig. (8)]. In [Fig. (9)], we show another




Fig. 10. Segmenting lesion in ultrasound breast image. (a) The original ultrasound image, with the initial model drawn on top, (b) The shape image based on edge map of the image, (c) The texture likelihood map, (d) The final segmentation result.




Fig. 11. Boundary finding in the pepper image. (a) Original image, with initial model drawn in blue. (b) The shape image derived from edge map, with edges drawn in yellow. (c) The intensity probability map derived based on model interior statistics. (d) Region of Interest (ROI) extracted. (e) Final segmentation result.

example in which we segment the Endocardium of the left ventricle in a noisy tagged MR heart image. Note that, due to the tagging lines and intensity inhomogeneity, the detected edges of the object are fragmented, and there are spurious small edges inside the region. In this case, the integration of both shape and texture information is critical in helping the model out of local minima. In [Fig. (10)], a metamorph model is used to extract the boundary of a lesion in an ultrasound image of the breast. On natural images, we show an example using the pepper image in [Fig. (11)]. Starting from a small model initialized inside the object, the model quickly deforms to the object boundary.

The Metamorph model evolution is computationally efficient, due to our use of the nonparametric texture representation and FFD parameterization of the model deformations. For all the examples shown, the segmentation process takes less than $200ms$ to converge on a 2Ghz PC station.

4 Conclusions

In this chapter, we have reviewed traditional shape-based deformable models, and introduced new frameworks that integrate region texture information into deformable models.

The new class of deformable models we proposed, Metamorphs, possess both boundary shape and interior intensity statistics. In Metamorphs the boundary

and region information are intergated within a common variational framework to compute the deformations of the model towards the correct object boundaries. There is no need to learn statistical shape and appearance models *a priori*. In our formulation, the model deformations are constrained so that the interior model statistics as it deforms remain consistent with the statistics learned from the past evolution of the model's interior. This framework represents a generalization of previous *parametric* and *geometric* deformable models, by exploiting the best features of both worlds. Segmentation using Metamorph models can be straightforwardly applied in 3D, and can handle efficiently the merging of multiple models that are evolving simultaneously.

Acknowledgement

This research has been funded by an NSF-ITR-0205671 grant to the first author.

References

1. A. A. Amini, Y. Chen, M. Elayyadi, and P. Radeva. Tag surface reconstruction and tracking of myocardial beads from SPAMM-MRI with parametric b-spline surfaces. *IEEE Transactions on Medical Imaging*, 20(2):94–103, 2001.
2. E. Bardinet, L. D. Cohen, and N. Ayache. A parametric deformable model to fit unstructured 3D data. *Computer Vision and Image Understanding*, 71(1):39–54, 1998.
3. V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. In *IEEE Int'l Conf. on Computer Vision*, pages 694–699, 1995.
4. T. Chan and L. Vese. Active contours without edges. *IEEE Trans. on Image Processing*, 10(2):266–277, 2001.
5. T. Chen and D. Metaxas. Gibbs prior models, marching cubes, and deformable models: a hybrid segmentation framework for medical images. In *Medical Imaging Computing and Computer-Assisted Intervention*, volume 2, pages 703–710, 2003.
6. L. D. Cohen and I. Cohen. Finite-element methods for active contour models and balloons for 2-D and 3-D images. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 15:1131–1147, 1993.
7. D. Decarlo and D. Metaxas. Blended deformable models. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 18(4):443–448, 1996.
8. J. M. Hammersley and P. Clifford. *Markov fields on finite graphs and lattices*. Preprint University of California, Berkeley, 1971.
9. X. Huang, D. Metaxas, and T. Chen. Metamorphs: Deformable shape and texture models. In *IEEE Conf. on Computer Vision and Pattern Recognition*, volume 1, pages 496–503, 2004.
10. X. Huang, N. Paragios, and D. Metaxas. Establishing local correspondences towards compact representations of anatomical structures. In *Proc. of Int'l Conf. on Medical Imaging Computing and Computer-Assisted Intervention, LNCS 2879*, pages 926–934, 2003.
11. M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *Int'l Journal of Computer Vision*, 1:321–331, 1987.
12. R. Malladi, J. Sethian, and B.C. Vemuri. Shape modeling with front propagation: A level set approach. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 17(2):158–175.

13. T. McInerney and D. Terzopoulos. A dynamic finite element surface model for segmentation and tracking in multidimensional medical images with application to cardiac 4d image analysis. *Computerized Medical Imaging and Graphics*, 19(1), 1995.
14. D. Metaxas. *Physics-Based Deformable Models*. Kluwer Academic Publishers, 1996.
15. D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Communications on Pure and Applied Mathematics*, 42(5):577–685, 1989.
16. S. Osher and N. Paragios (Editors). *Geometric Level Set Method in Imaging, Vision and Graphics*. Springer Verlag, 2002.
17. N. Paragios and R. Deriche. Geodesic active regions and level set methods for supervised texture segmentation. *Int'l Journal of Computer Vision*, 46(3):223–247, 2002.
18. J. Park, D. Metaxas, A. Young, and L. Axel. Deformable models with parameter functions for cardiac motion analysis. *IEEE Trans. on Medical Imaging*, 15(3), 1996.
19. R. Ronfard. Region-based strategies for active contour models. *International Journal of Computer Vision*, 13(2):229–251, 1994.
20. T. W. Sederberg and S. R. Parry. Free-form deformation of solid geometric models. In *Proceedings of the 13th Annual Conference on Computer Graphics*, pages 151–160, 1986.
21. J. A. Sethian. *Level Set Methods: Evolving interface in Geometry, Fluid Mechanics, Computer Vision and Material Science(First Edition)*. Cambridge University Press, 1996.
22. L. H. Staib and J. S. Duncan. Boundary finding with parametrically deformable models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(11):1061–1075, 1992.
23. A. Tsai, A. Willsky, and A. J. Yezzi. A statistical approach to snakes for bimodal and trimodal imagery. In *IEEE Int'l Conf. on Computer Vision*, volume 2, pages 898–903, 1999.
24. L. A. Vese and T. F. Chan. A multiphase level set framework for image segmentation using the Mumford and Shah model. *Int'l Journal of Computer Vision*, 50(3):271–293, 2002.
25. R. Whitaker. Volumetric deformable models. In *the Third Conference on Visualization in Biomedical Computing (VBC'94)*, 1994.
26. M. Worring, A. W. M. Smeulders, L. H. Staib, and J. S. Duncan. Parameterized feasible boundaries in gradient vector fields. *Computer Vision and Image Understanding*, 63(1):135–144, 1996.
27. C. Xu, D. L. Pham, M. E. Rettmann, D. N. Yu, and J. L. Prince. Reconstruction of the human cerebral cortex from magnetic resonance images.
28. C. Xu and J. L. Prince. Generalized gradient vector flow external forces for active contours. *Signal Processing: An International Journal*, 71(2):131–139, 1998.
29. A. Yezzi, S. Kichenessamy, A. Kumar, P. Olver, and A. Tannebaum. A geometric snake model for segmentation of medical imagery. *IEEE Trans. on Medical Imaging*, 16(2):199–209, 1997.
30. S. Zhu and A. Yuille. Region Competition: Unifying snakes, region growing, and Bayes/MDL for multi-band image segmentation. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 18(9):884–900, 1996.