

Integrated Surface Model Optimization for Freehand Three-Dimensional Echocardiography

Mingzhou Song*, Robert M. Haralick, Florence H. Sheehan, and Richard K. Johnson

Abstract—The major obstacle of three-dimensional (3-D) echocardiography is that the ultrasound image quality is too low to reliably detect features locally. Almost all available surface-finding algorithms depend on decent quality boundaries to get satisfactory surface models. We formulate the surface model optimization problem in a Bayesian framework, such that the inference made about a surface model is based on the integration of both the low-level image evidence and the high-level prior shape knowledge through a pixel class prediction mechanism. We model the probability of pixel classes instead of making explicit decisions about them. Therefore, we avoid the unreliable edge detection or image segmentation problem and the pixel correspondence problem. An optimal surface model best explains the observed images such that the posterior probability of the surface model for the observed images is maximized. The pixel feature vector as the image evidence includes several parameters such as the smoothed grayscale value and the minimal second directional derivative. Statistically, we describe the feature vector by the pixel appearance probability model obtained by a nonparametric optimal quantization technique. Qualitatively, we display the imaging plane intersections of the optimized surface models together with those of the ground-truth surfaces reconstructed from manual delineations. Quantitatively, we measure the projection distance error between the optimized and the ground-truth surfaces. In our experiment, we use 20 studies to obtain the probability models offline. The prior shape knowledge is represented by a catalog of 86 left ventricle surface models. In another set of 25 test studies, the average epicardial and endocardial surface projection distance errors are 3.2 ± 0.85 mm and 2.6 ± 0.78 mm, respectively.

Index Terms—Echocardiography, image analysis, image shape analysis, surface reconstruction.

I. INTRODUCTION

ULTRASOUND imaging is done in real time, making it more suitable for cardiology than other imaging modalities such as magnetic resonance imaging. It is preferable also because the equipment is portable, safe, and noninvasive. Our goal is to create a three-dimensional (3-D) left ventricle (LV) surface model, including the *epicardium* (EPI), the outer surface of the LV, and the *endocardium* (ENDO), the inner surface of the

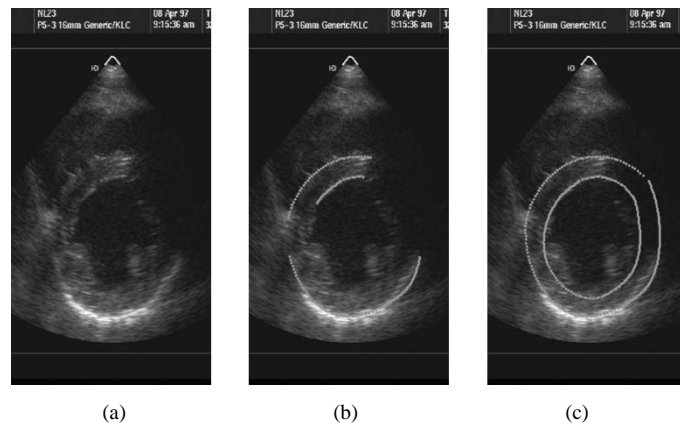


Fig. 1. An ultrasound image of the LV and surface model contours. (a) Short axis view. (b) Visible contour overlaid. (c) Complete contour overlaid.

LV, from multiple-view two-dimensional (2-D) ultrasound images of the heart at *end systole* or *end diastole*. During image acquisition, a sonographer chooses multiple appropriate positions for the transducer; the 3-D spatial location and orientation of the corresponding imaging planes are quantitatively obtained using a magnetic field tracking device [1]. Imaging planes from different views may not be parallel and can intersect each other at quite arbitrary angles. A fairly good-quality ultrasound image of the heart is shown in Fig. 1(a). The low image quality is due to artifacts such as strong speckle noises, image dropout, low image resolution, and fast moving small structures of the heart.

When examining an echocardiographic image, doctors or sonographers register the image with a mental picture of a generic LV. If they want to do quantitative analysis of the LV for clinical research, they will draw some borders on the images where they think the endocardial or epicardial boundaries are most likely to occur, according to both the image and their experience of the LV shape. They can further build a specific 3-D surface model of the LV with these borders.

Many previous approaches to automated surface modeling have used a two-stage or an iterative two-stage strategy: isolated feature detection and model fitting stages. They do not mimic the human reasoning process in guiding the low-level feature detection with the prior shape knowledge. The suboptimality manifests when the image quality is too poor to reliably detect features. On sparse and noisy ultrasound images, the prior shape knowledge of the LV must be considered during surface model optimization.

We apply an integrated surface model optimization approach to 3-D echocardiography. Our approach is a unified framework that is integrated in one step. Accordingly, there is neither image

Manuscript received December 2, 2001; revised July 6, 2002. Asterisk indicates corresponding author.

*M. Song is with the Department of Computer Science, Queens College, City University of New York, Flushing, NY 11367 USA (e-mail: msong@cs.qc.edu).

R. M. Haralick is with the Doctoral Program in Computer Science, Graduate Center of City University of New York, New York, NY 10016 USA (e-mail: haralick@gc.cuny.edu).

F. H. Sheehan is with the Division of Cardiology, Department of Medicine, University of Washington, Seattle, WA 98195 USA (e-mail: sheehan@u.washington.edu).

R. K. Johnson is with Quantographics, Inc., Mercer Island, WA 98040 USA (e-mail: richj@quantographics.com).

Digital Object Identifier 10.1109/TMI.2002.804433

segmentation nor edge detection¹ based on local image operation. We assign each pixel a class probability profile instead of an exclusive class label. We prove the optimality of the integrated approach by a surface inference theorem. The integrated approach seeks a 3-D surface model that has the greatest posterior probability for the observed images. The posterior probability of a surface model is in proportion to its prior probability and its probability of producing the given images.

A catalog of known surface models and other user input information determines the prior probability of a surface model. The catalog represents admissible shapes and generates new surface models by combining its members. The surface appearance probability is the summation of the product between the pixel appearance and pixel class prediction probabilities. We model the pixel appearance probability by a nonparametric optimal quantization technique [2]. The pixel class prediction process has a physical simulation part and a pixel class prediction probability modeling part. We obtain the pixel appearance and class prediction probability models offline using a generalized expectation maximization algorithm.

Our integrated surface optimization approach is designed to handle problems where the image quality is low and the prior shape knowledge is important, as in 3-D echocardiography. We performed an experiment on 20 training and 25 test clinical studies. The results on the test set showed an average epicardial surface projection distance error of 3.2 ± 0.85 mm and an average endocardial surface projection distance error of 2.6 ± 0.78 mm. The error was measured between the optimized and the ground-truth surfaces.

This paper is organized as follows. We review the related work of surface model optimization in Section II. We discuss the integrated approach in Section III. We describe how we represent the LV and its prior shape knowledge in Section IV. We give the algorithm for joint estimation of the pixel appearance and class prediction probability models in Section V. We present the LV surface model optimization algorithms in Section VI. We report the LV optimization results in Section VII. We conclude the integrated approach in Section VIII.

II. RELATED WORK

Previous methods on surface reconstruction include those for noise-free or very low-noise image data and those for sparse and noisy image data. We treat manually processed data as low-noise data. Frangi *et al.* [3] provide an extensive review on 3-D cardiac modeling.

A. Surface Modeling From Low-Noise Data

Hoppe *et al.* [4]–[6] report a surface reconstruction method for dense and low-noise data. Their method creates an initial triangular mesh model from 3-D unorganized points directly and then optimizes the mesh model by its goodness of fit to the data and the local smoothness. The final mesh model is a piecewise smooth representation of the target surface based on subdivision. Legget *et al.* [1] modify Hoppe's method and apply it in 3-D reconstruction of the LV from multiple-view 2-D ultrasound heart images. They build a prototype mesh model offline

¹In accordance with usage in the image-processing literature, edge detection means making a definite yes/no decision as to whether a pixel is an edge.

for the LV surface. The data are manually labeled 3-D points. The prototype mesh model is freely deformed to fit the 3-D points, sparser than the data in Hoppe's experiments. The prior shape knowledge is employed in the offline model creation. So the application of the prior shape knowledge is weak, and the accuracy of the optimization relies on the quality of the manually extracted and labeled 3-D points.

B. Surface Modeling From Sparse and Very Noisy Data

When the data are sparse and very noisy, prior shape knowledge comes into play. The local smoothness can often be modeled using splines. Kass *et al.* [7] propose a 2-D snakes model. Chalana *et al.* [8] sweep the 2-D snakes model along the time axis to track the 2-D contour of LV in a sequence of images.

Due to their simplicity, there is a wide range of work on parametric shape models. In 2-D, ellipses or elliptical arcs are the most used shape model [9]–[11]. In 3-D, the cylinder, ellipsoid, disc, or some of their combinations are the most commonly adopted shape models. Brinkley [12] slightly deforms spheres to model organs. Stetten *et al.* [13], [14] model the LV in 3-D with a combination model of a half-sphere, a cylinder, and a slab. The more complex superquadrics model is proposed by Pentland [15], with which Solina and Bajcsy [16] approximate 3-D objects and Bardin *et al.* [17] and Park *et al.* [18] model the LV.

Since these parametric models are not complex enough to capture the shape concept of all natural objects, further free-form deformation of these models after their initial fit to the data might be necessary. However, the free-form deformation might produce inadmissible shapes, because either there is no constraint on the free form deformation or the constraints cannot be summarized systematically.

Statistical shape models capture both shape complexity and variations, avoiding the variance of the unconstrained free-form deformation. They can represent any plausible shape within the constraints imposed by the variability in the training set. Based on Brinkley's [19] radial contour model, Altman and Brinkley [20] use a multivariate normal distribution to model the radial distances from a center point to the contour. Cootes *et al.* [21], [22] first suggest the 2-D active shape model, representing a shape with a vector composed by a set of landmark points and interpolated points. They capture the shape variability by the first several principal components of the shape vectors. They have applied active shape models in face recognition, finger finding, and 2-D LV contour detection from ultrasound images. They extend their statistical shape modeling work in [23]. Blake and Isard [24], [25] design the 2-D active contour models, which are similar to the active shape models. They do not use the landmark points directly to represent a shape. Instead, they fit splines to the 2-D shape contour and use the control points of the splines to represent the shape, so that the shape vector is in the spline space. Hence, their model incorporates both the local smoothness constraint and the global shape information. Jacob *et al.* [26] apply the active contour model in tracking 2-D heart contours in echocardiographic sequences. Extending the 2-D statistical shape models of Cootes *et al.* into 3-D, Counce and Taylor [27] create a 3-D point distribution model to represent cortical sulci. Lorenz and Krahnstöver [28] attempt the

point-based 3-D statistical shape models for anatomical objects. A difficulty in these 3-D point-based approaches is the point correspondence problem. Different approaches to solving the point correspondence problem are summarized in [27]. Hwang *et al.* [29] investigate a distance transform neural-network-based 3-D left ventricle recognition framework. Statistical shape-modeling methods based on landmarks are proposed in [30] and [31], but they are mainly concerned with shape analysis rather than surface reconstruction.

C. Fitting Shape Models to Image Data

Usually the data given for object reconstruction are not points on the object surface, but images of the object. Naturally, pixels on the surface are first detected from the images. The most used feature detectors are gradient-based edge detectors, e.g., Canny's edge detector [32] and the facet model edge detector [33]. A shape model is then fitted to the detected points. In this approach, the feature detection and the model fitting are separated. The shape knowledge has no impact on the feature detection phase. Coppini *et al.* [34] describe a neural-network-based two-stage approach. Chakraborty *et al.* [35] integrate the gradient and region information for pixel classification. The region information is modeled by Markov random fields with no shape statistics applied.

Methods to reconstruct a particular shape model from low-quality images with the guidance of shape statistics have not made much progress beyond the two-stage approach of feature detection and model fitting. The framework for snakes or active contour models does blend the two stages without feature detection, but it lacks a systematic way of training because they are not formulated in a probabilistic framework.

Cootes *et al.* [21], [22] model image pixels by the statistical active appearance models. They combine the active appearance models and the active shape models to find the best 2-D contour from images. Pixel classification decisions are still made but might change during the model fitting iterations. This iterative classification scheme allows shape statistics to guide the local feature detection. Montagnat *et al.* [36] deform their 3-D mesh model to nearby voxels that are most likely to be on the boundary. As far as we have found in the literature, only Mignotte and Meunier [37] have explored the idea of 2-D cardiac contour tracking from images in a probabilistic framework without an explicit feature detection stage.

III. THE INTEGRATED APPROACH TO SURFACE MODEL OPTIMIZATION

Given the observed images, what is the surface model that best explains these images? Owing to the noise in the imaging process, every possible solution has an associated uncertainty. If we assign every solution a probability, then the best solution has the maximum probability given the observed images. We use the Bayesian framework to formulate the surface model optimization problem, because it is the unique consistent induction framework under the Cox–Jaynes axioms [38], [39]

$$\max_{\Theta} p(\Theta|Z) = \max_{\Theta} \frac{p(\Theta)p(Z|\Theta)}{p(Z)} \propto \max_{\Theta} p(\Theta)p(Z|\Theta)$$

where Θ is the surface model, Z is the feature vector of a pixel, $p(\Theta|Z)$ is the posterior probability of Θ given Z , $p(\Theta)$ is the prior probability of surface model Θ , and $p(Z|\Theta)$ is the conditional probability of observed feature vector Z given a surface model Θ . We also call $p(Z|\Theta)$ the surface appearance model, representing the overall imaging system behavior.

The prior probability $p(\Theta)$ can be assessed in ways depending upon the application. For a surface model, it can be prior probability characterizing smoothness, or the shape of the objects, or simply some user input points. $p(Z)$ is the probability of observing a particular feature vector Z , the knowledge of which is only necessary when the exact posterior probability is desired. Although theoretically $p(Z|\Theta)$ can be computed, the task is difficult because of the many degrees of freedom of the observed feature vectors.

To avoid the direct computation of $p(Z|\Theta)$, the degrees of freedom of the feature vector Z can be reduced by assigning each pixel a class label, denoting whether the pixel is visibly located on a certain surface. This process is known as feature detection. We use Y to denote the class label of a pixel. If we can detect the class labels for all the pixels, we can search for a surface model Θ^* that best fits the class labels, instead of best fitting the original images. These two stages—feature detection and model fitting—abstract almost all previous work, which is summarized as Algorithm 1.² $P(Y|Z)$ is the posterior of a class label Y given the feature vector Z . $p(Z|Y)$ is the likelihood of the class label Y for the feature vector Z . $p(\Theta|Y)$ is the posterior of the model surface Θ given the class label Y . $P(Y|\Theta)$ is the likelihood of the surface model Θ for the class label Y . Direct computation of $p(Z|\Theta)$ is avoided here. Instead, estimates of $p(Z|Y)$ and $P(Y|\Theta)$ are needed, as well as the priors $P(Y)$ and $p(\Theta)$. This two-stage approach does not fall into an overall Bayesian framework. If we can detect the class label Y from the feature vector Z with good confidence, the two-stage approach will work. Otherwise, if we have to detect features on fuzzy images, the two-stage framework does not yield an optimal surface model because the detected class label Y^* is unreliable.

Algorithm 1 Two-stage surface model optimization.

Stage 1. Feature detection. Find Y^* that solves

$$\max_Y P(Y|Z) \propto \max_Y P(Y)p(Z|Y).$$

Stage 2. Model fitting. Find Θ^* that solves

$$\max_{\Theta} p(\Theta|Y^*) \propto \max_{\Theta} p(\Theta)P(Y^*|\Theta).$$

The integrated approach that we have developed avoids making exclusive class label decisions for each pixel by

²We use $p(\cdot)$ to represent probability density function and $P(\cdot)$ to represent probability mass function.

integrating the two stages into one. We will not assign a class label to each pixel, but will profile each pixel by probabilities of having different class labels.

A. The Integrated Approach

In the integrated framework, we still introduce the class label Y to avoid the direct offline estimation and online computation of $p(Z|\Theta)$. A class label is not actually directly observable, but serves as a hidden bridge between the feature vector Z and the surface model Θ , by marking each pixel to be either on or off a certain surface model.

Before we proceed, consider a scenario of surface reconstruction. If somehow the class label of each pixel is already known, the feature vector would not provide more information in the inference about the surface model Θ . Therefore, we make the following assumption:

$$p(\Theta|Z, Y) = p(\Theta|Y) \quad (\neq p(\Theta|Z)) \quad (1)$$

implying that the surface model Θ and the feature vector Z are conditional independent given the class label Y .

Theorem 1—Integrated Surface Inference: With the assumption $p(\Theta|Z, Y) = p(\Theta|Y)$, the posterior probability of a surface model Θ given the observed feature vector Z can be written as the summation form

$$p(\Theta|Z) = \frac{p(\Theta)}{p(Z)} \sum_{y=1}^K P(Y = y|\Theta)p(Z|Y = y) \quad (2)$$

where K is the total number of classes.

Proof: By properties of probability and assumption

$$\begin{aligned} p(\Theta|Z) &= \sum_y p(\Theta, Y = y|Z) \\ &= \sum_y p(\Theta|Y = y, Z)P(Y = y|Z) \\ &= \sum_y p(\Theta|Y = y) \frac{p(Z|Y = y)P(Y = y)}{p(Z)} \\ &= \frac{p(\Theta)}{p(Z)} \sum_y P(Y = y|\Theta)p(Z|Y = y). \end{aligned} \quad (3)$$

The integrated surface inference theorem leads to the *integrated approach of surface model optimization*

$$\Theta^* = \arg \max_{\Theta} p(\Theta) \sum_y P(Y = y|\Theta)p(Z|Y = y) \quad (4)$$

where Θ^* is an optimal surface model. Evidently from the second line in (3), the two-stage approach is a special case of the integrated approach if $P(Y|Z)$ is an impulse function of Y .

We call the conditional probability mass function $P(Y|\Theta)$ the *pixel class prediction* (PicPre) probability model, emphasizing the physical meaning that the class label Y can be predicted probabilistically from a given surface model Θ . We call the con-

ditional probability density function $p(Z|Y)$ the *pixel appearance* (PixApp) probability model because $p(Z|Y)$ depicts probabilistically the appearance of a pixel for the class label Y .

The interpretation behind the integrated approach is as follows. An image pixel with feature vector Z is assigned a likelihood profile $p(Z|Y)$ of being from different classes Y . Another class probability $P(Y|\Theta)$ profile is predicted from a model. When the likelihood profile $P(Y|Z)$ and the predicted class probability profile $P(Y|\Theta)$ match well, the surface model that generates the predicted class probability profile is a good explanation of the images.

The work of Mignotte and Meunier [37] follows the principle of the integrated approach. However, they explain their idea intuitively without deriving (4), and their formulation uses a prescribed simple distribution for $P(Y|\Theta)$.

B. The Pixel Appearance Probability Model

The appearance of a pixel is defined by its local information, such as intensity, contrast, directional derivatives, etc. The appearance of a pixel is a result of the structure type at the pixel location and the imaging process; it is not strictly a function of the structure types. Fig. 1(a) shows the ultrasound image of the LV, and Fig. 1(b) shows the same image with the visible boundary overlaid. Fig. 1(c) overlays the complete contour of the underlying surface model on the original image. It is quite obvious that the pixels on the underlying surface contour do not have uniform appearance everywhere: some pixels are bright with high contrast, while others do not differ too much with the background. The background pixels also have variable appearance.

In ultrasound images, signals arriving at an interface between media with different acoustic impedance produce strong echo when the angle of incidence is near perpendicular; signals arriving at an interface at near tangential angles produce very weak echo. Thus the image intensity and its spatial variation are important. We fit a cubic facet model [33] to the pixel intensities in a local window centered at each pixel. With the facet, a surface patch in 3-D, we are able to analytically derive all first- and second-order derivatives. The derivatives carry spatial variation information. Based on these considerations, each pixel feature vector contains³:

- 1) the spatially and temporally smoothed pixel intensity value, capturing absolute strength of echo signals;
- 2) the third-order facet model approximation of the pixel intensity value, using a larger spatial window than feature 1), giving the absolute strength of echo signals but on a larger scale;
- 3) the directional derivative along the gradient direction, providing a local measure of edginess;
- 4) the minimum second directional derivative, among second derivatives along all directions, indicating relative strength of echo signals;
- 5) the directional derivative from a point inside the LV, to help distinguish ENDO and EPI surfaces. The inner point is derived from user input points.

³In this paper, we concentrate on integrated optimization. The selected features work well, but it is possible that other features might produce even better results with our approach.

For feature 1), we used a 3-D median filter with a window of 5 pixels \times 5 pixels \times 5 frames. For features 2) to 5), we used a facet window of 21 pixels \times 21 pixels. For feature 5), we used a 3-D median filter with a window of 17 pixels \times 17 pixels \times 17 frames. Examples of the feature vector maps are shown in Fig. 2.

The PixApp probability model describes the pixel appearance uncertainty for a given pixel class. We design a nonparametric optimal quantization technique [2] to characterize the pixel appearance probability model $p(Z|Y)$. Nonparametric models are known to have less modeling biases than parametric models. Nonparametric models typically require large sample size to reduce both the bias and the variance. Since the model $p(Z|Y)$ is pixel based, it is not an issue to have enough samples. But the major concern is CPU cycles. Standard nonparametric approaches, such as kernel methods, store all samples and run in quadratic time in sample size. For a multivariate sample of size on the order of millions—an easy thing for image pixels—the online computation of kernel methods is formidable. We use the much more efficient grid model to represent $p(Z|Y)$. The appearance space is optimally quantized into cells in the following senses: 1) the partition of the space is compact and statistically efficient by adapting to the data; 2) the density estimate of each cell is optimally smoothed to control its variance.

C. Pixel Class Prediction

Pixel class prediction associates every pixel on an imaging plane with some physical properties of a given surface model and its environment. Pixel class prediction and classification are fundamentally different. We denote the deterministic prediction by $Y|\Theta$, where each pixel has an exclusive class assignment. We represent the probabilistic prediction by the conditional probability $P(Y|\Theta)$. The class probability profile provides a soft class prediction, allowing a more precise relation between a pixel and its predicted class to be captured. When a 3-D surface model maps to a 2-D image, two phenomena occur: 1) a 3-D point is transformed to a 2-D pixel location and 2) the physical properties of the 3-D point yield a 2-D pixel intensity. The physical simulation predicts images by functional modeling of the imaging system. In [40] and [41], we have implemented an ultrasound imaging simulation system to synthesize echocardiographic images from a 3-D surface model of the LV. The simulation software is capable of performing backscattering, attenuation, and reflection, implemented by a ray-tracing algorithm. In this paper, we only simulate reflections, since our purpose is to predict the systematic image dropout rather than the stochastic behavior of the speckle noise. The dropout is mostly due to weak reflection at interfaces. The randomness is accounted for in the PixPre and PixApp probability models.

Let classes 1 to $K-1$ correspond to the surfaces in the surface model. An additional class K labels the background. For an LV surface model, $K = 3$, and 1, 2, and 3 correspond to EPI, ENDO, and background respectively. The distance from a pixel p to its closest class- y neighbor pixel q on the simulated image is denoted by $d(y)$. The intensity of the neighbor pixel q is denoted by $I(y)$. The overall probabilistic prediction is shown in Fig. 3. A physical simulation produces simulated images from a surface model Θ . $d(y)$ and $I(y)$ of each pixel location in a sim-




Fig. 2. Feature maps extracted from the parasternal long-axis view. (a) The original image; (b)–(f) maps of the five features.

ulated image can be found by the distance transform [33], [42] efficiently. Then the pixel class probability profile $P(Y|\Theta)$ is predicted by probability modeling.

During the imaging process, a point on the surface may be transformed to a pixel looking more like the background; a point off the surface may be transformed to a pixel as if on the surface. Instead of predicting that a pixel coming from a point on surface y must have label y , the probabilistic model allows variation. An off-surface 3-D point closer to an on-surface 3-D point may appear as a 2-D pixel with similar location and intensity with the type of the 2-D on-surface pixel. An off-surface point that has a stronger on-surface neighbor point is more likely to appear as a pixel that is similar to the type of the neighbor on-surface pixel. Therefore, we use the following intensity exponential decay model:

$$P(Y = y|\Theta) = \begin{cases} \frac{I(y)e^{-\lambda_y d(y)}}{\beta + \sum_{k=1}^{K-1} I(k)e^{-\lambda_k d(k)}}, & y = 1, 2, \dots, K-1 \\ \frac{\beta}{\beta + \sum_{k=1}^{K-1} I(k)e^{-\lambda_k d(k)}}, & y = K \end{cases} \quad (5)$$

where λ is a nonnegative parameter that controls the rate of decay, $\lambda_1, \lambda_2, \dots, \lambda_{K-1}$ control the rate of decay for different classes, and β is a nonnegative parameter that corresponds to the strength of a pixel's being the background. $d(k)$ is the distance from the pixel to its closest class k neighbor pixel.




Fig. 3. Probabilistic pixel class prediction.

IV. LEFT VENTRICLE SURFACE MODELS AND THE PRIOR SHAPE KNOWLEDGE

Epicardial and endocardial surfaces are each represented by a triangular mesh model, which is capable of capturing irregular shapes better than a parametric model. A triangular mesh model of the LV is shown in Fig. 4.

We model the prior probability of mesh surface models by a hybrid approach. A catalog of 86 LV models was established by reconstructions from the manual delineations on ultrasound images. The members in the catalog are representatives of possible surface models and span a subspace in the full surface model space. New models not in the catalog are approximated by convex combinations of the surface models in the catalog. The catalog is a geometric representation of the prior knowledge about surface model shapes. The vertices of the mesh models are ordered according to their anatomical position.

An explicit prior probability can be formed by some user input points. There are two types: landmark points and surface points. Landmark points are usually special geometric points that can be identified with small variance. Landmark points may or may not be on a surface. For example, the center of the mitral valve is a landmark point. For the LV models in the catalog developed for this study, every surface model has the same set of landmarks, and there is a one-to-one relation between landmarks of one surface model and those of another. Surface points are located on surfaces, but there is no correspondence between surface points on one surface and those on another. For example, any point actually located on the ENDO surface is a surface point. In our work, the user inputs three landmark points: the apex on ENDO, centers of the mitral and aortic valves, and four ENDO surface points, approximately equally distributed around the LV contour in the short-axis view.

Owing to image quality and human factors, user input points have uncertainty in themselves. User input landmark points may deviate from the true positions randomly. The surface points may not sit exactly on the true surfaces. For the known surface models, we gather the distance error statistics for these two types of user input points and create parametric probability models to represent the uncertainty.

V. PIXAPP AND PICPRE PROBABILITY MODEL ESTIMATION

Since the class label Y is not directly observable, optimal estimation of the PicPre probability model has to interact with the estimation of the PixApp probability model in the integrated framework. To get the overall optimal probability models, we solve the simultaneous estimation of the PixApp and PicPre probability models by a generalized expectation maximization (EM) algorithm.




Fig. 4. A triangular mesh surface model of the LV.

A. Joint Estimation of PixApp and PicPre Probability Models

In the offline training phase, the overall goal is to make an accurate and consistent estimation of $p(Z|\Theta)$. To understand the probability estimation in the integrated framework, it is necessary to discuss briefly the estimation problem in the two-stage approach. In the feature detection of the two-stage approach, the class conditional probability of feature given class $p(Z|Y)$ is used. For a given Z , the label Y maximizing $p(Z|Y)$ is found. However, the class label is not observed data. In the surface reconstruction stage, the PicPre probability model $P(Y|\Theta)$ is used. Using the class labels and known surface models, $P(Y|\Theta)$ can be estimated. The only problem with these two-stage estimations is that the uncertainty of class label Y as described by $P(Y|\Theta)$ is not taken into account in the estimation of $p(Z|Y)$, which can degrade the performance seriously when the uncertainty of class labels for given surface models is prominent.

In the integrated framework, we still need to estimate the PicPre $P(Y|\Theta)$ and the PixApp $p(Z|Y)$ probability models. However, we do not have to make a decision on the class label Y of each pixel because every possibility of Y is considered. Since we have decided that $P(Y|\Theta)$ is a parametric model and

$p(Z|Y)$ is a nonparametric model, the overall model $p(Z|\Theta)$ is a hybrid model. On one hand, maximum likelihood estimation for $p(Z|\Theta)$ requires joint estimation of the PixApp and the PicPre probability models. On the other hand, joint estimation of a hybrid model poses a computational challenge.

Although it is typically a solution to maximum-likelihood parameter estimation with missing or hidden variables, the EM algorithm performs on densities that can be written in an integral form. The EM algorithm has long been used in medical imaging [43], [44]. The missing or hidden variable is precisely the integral variable. Whether the target density is parametric, nonparametric, or hybrid will affect neither the applicability nor the convergence of the EM algorithm. In the integrated model, the goal is to maximize $\mathbf{E}[\log p(Z|\Theta)]$ over all possible probability densities $p(Z|\Theta)$ (not over all possible Θ). When $p(Z|\Theta)$ is written as the integrated form, Y is the missing or hidden variable. Instead of maximizing $\mathbf{E}[\log p(Z|\Theta)]$, the EM algorithm maximizes an approximation of $\mathbf{E}[\log p(Y, Z|\Theta)]$ over $p(Y, Z|\Theta)$ in its iterations. The sense of the EM algorithm lies in that maximization of $\mathbf{E}[\log p(Y, Z|\Theta)]$ is substantially computationally easier than that of $\mathbf{E}[\log p(Z|\Theta)]$.

Evidently, $p(\Theta|Z, Y) = p(\Theta|Y)$ implies $p(Z|Y, \Theta) = p(Z|Y)$. Therefore, $p(Y, Z|\Theta) = P(Y|\Theta)p(Z|Y, \Theta) = P(Y|\Theta)p(Z|Y)$. Hence

$$\begin{aligned} & \max_{p(Y, Z|\Theta)} \mathbf{E}_{\pi_m} [\log p(Y, Z|\Theta)] \\ &= \max_{p(Y, Z|\Theta)} \mathbf{E}_{\pi_m} [\log P(Y|\Theta)] + \mathbf{E}_{\pi_m} [\log p(Z|Y)] \\ &= \max_{P(Y|\Theta)} \mathbf{E}_{\pi_m} [\log P(Y|\Theta)] + \max_{p(Z|Y)} \mathbf{E}_{\pi_m} [\log p(Z|Y)] \end{aligned} \quad (6)$$

which is the maximization step of the EM algorithm. The M-step has been separated into two independent optimization problems. One is the parametric estimation of the PicPre probability model, and the other is the nonparametric estimation of the PixApp probability model. Replacing $p(Y, Z|\Theta)$ by $P(Y|\Theta)p(Z|Y)$, we give Algorithm 2.

Algorithm 2 Estimate-Integrated-Model

Initialization:

$$P_0(Y|\Theta) \text{ and } p_0(Z|Y).$$

Iteration:

1) E-step

$$\pi_m(Y|Z, \Theta) = \frac{P_m(Y|\Theta)p_m(Z|Y)}{\sum_k P_m(Y=k|\Theta)p_m(Z|Y=k)} \quad (7)$$

$$\phi_m(P(Y|\Theta)) = \mathbf{E}_{\pi_m} [\log P(Y|\Theta)] \quad (8)$$

$$\psi_m(p(Z|Y)) = \mathbf{E}_{\pi_m} [\log p(Z|Y)]. \quad (9)$$

2) M-step

$$P_{m+1}(Y|\Theta) = \arg \max_{P(Y|\Theta)} \phi_m(P(Y|\Theta)) \quad (10)$$

$$p_{m+1}(Z|Y) = \arg \max_{p(Z|Y)} \psi_m(p(Z|Y)). \quad (11)$$

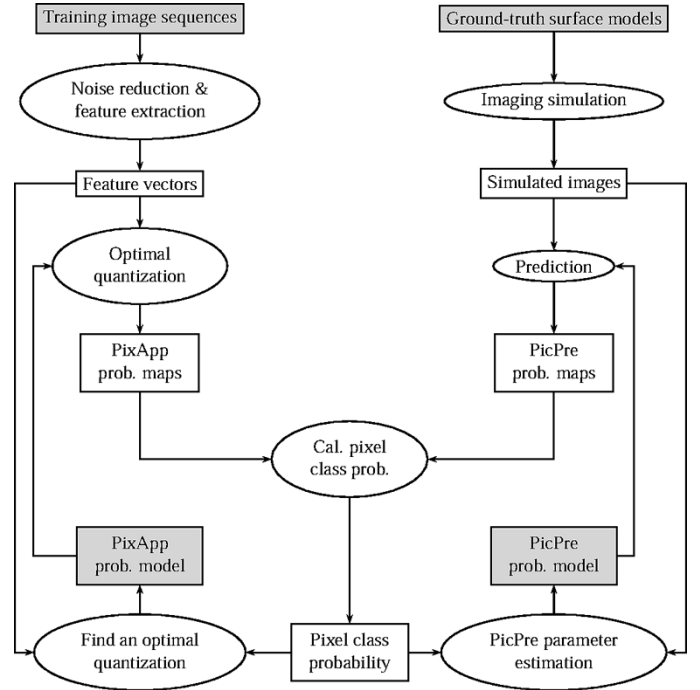


Fig. 5. Estimation of PixApp and PicPre probability models.

When the M-step returns a suboptimal solution that does not decrease $\phi_m(p(Y, Z|\Theta))$, the algorithm is known as a generalized EM algorithm. It is shown in [45] that both the original and the generalized EM algorithms increase the targeted expected log likelihood $\mathbf{E}[\log p(Z|\Theta)]$ monotonically.

The Estimate-Integrated-Model algorithm differs from the two-stage estimation solution in (7). In the two-stage approach, every pixel is assigned a unique class label y , equivalent to setting $\pi_m(Y|Z, \Theta) = \delta(Y - y)$. Here, $\pi_m(Y|Z, \Theta)$ signifies the probability profile of class labels for given images and surface model. In addition, the Estimate-Integrated-Model algorithm iterates across the two steps, while the two-stage approach does them only once. Fig. 5 gives a more concrete description of Algorithm 2 to explain how to obtain the PixApp and the PicPre probability models.

B. PixApp Probability Model Estimation

One of the two expectations to be maximized in the M-step is $\mathbf{E}_{\pi}[\log p(Z|Y)]$. The expectation is with respect to both Y and Z . However, the unknown conditional probability $P(Y|Z, \Theta)$ is replaced by an approximation $\pi(Y|Z, \Theta)$. Hence $\mathbf{E}_{\pi}[\log p(Z|Y)]$ can be written as

$$\mathbf{E}_{\pi} [\log p(Z|Y)] = \int p(z|\Theta) \sum_{k=1}^K \pi(Y=k|z, \Theta) \cdot \log p(z|Y=k) dz. \quad (12)$$

By taking sample average log likelihood as the expected one

$$L = \frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \pi(y_n = k|z_n, \Theta) \log p(z_n|y_n = k). \quad (13)$$

Thus maximization of $\mathbf{E}_{\pi}[\log p(Z|Y)]$ is approximated by that of L . As Z is usually a continuous vector, $p(Z|Y)$ is a probability density function conditioned on the discrete variable Y .

We introduce an optimal quantization technique to find $p(Z|Y)$. Optimal quantization partitions the space into cells, such that they give a compact representation of the feature vector Z . The pattern formed by the cells maximizes a quantizer performance measure that is the summation of the weighted log likelihood and the entropy. The optimal quantization approach easily incorporates the class weights of each datum by allowing fractional sample size. Entropy insures the consistency of the density estimates. Otherwise, the estimated density would be a function of superimposed impulses, which will give the poorest performance on unseen data not in the training set. The final density estimate of each cell for a given class is a nonzero quantity obtained by smoothing the empirical density of each cell. We assign every cell with some nonzero density estimate by smoothing with a control parameter optimized by cross-validation. Although smoothing also maintains the consistency of the density estimates, its effect cannot be achieved by the mere usage of entropy. Once they are estimated, the density functions are stored in a grid. With a properly chosen cell shape, the grid can be retrieved much more efficiently than for other standard nonparametric methods. While it is true that nonparametric methods always benefit from larger sample size, there is an important property that is unique to optimal quantization that the performance can always improve with more quantization levels.

For a pixel in an ultrasound image of the heart, the original feature vector Z has five dimensions. We reduce Z to a 3-D vector X by taking the first three principal components of Z . We actually perform optimal quantization on X instead of Z . We show the one-dimensional (1-D) and 2-D marginal densities of the 3-D PixApp probability densities $p(X|Y)$ in Figs. 6 and 7, respectively. Classes $Y = 1, 2, 3$ represent EPI, ENDO, and background pixels, respectively. Fig. 8 shows three class PixApp probability maps of a given image, obtained with the estimated PixApp probability model.

C. PicPre Probability Model Estimation

The other expectation to be maximized in the M-step is $\mathbf{E}_\pi[\log P(Y|\Theta)]$. The expectation is on both Z and Y , where Z is implicitly expressed in $\pi(Y|Z, \Theta)$. $\mathbf{E}_\pi[\log P(Y|\Theta)]$ is an estimate of $\mathbf{E}[\log P(Y|\Theta)]$ with $P(Y|Z, \Theta)$ replaced by $\pi(Y|Z, \Theta)$. Therefore, we have

$$\mathbf{E}_\pi[\log P(Y|\Theta)] = \int p(z|\Theta) \sum_{k=1}^K \pi(Y = k|z, \Theta) \cdot \log P(Y = k|\Theta) \cdot dz \quad (14)$$

We can further obtain an estimate of $\mathbf{E}_\pi[\log P(Y|\Theta)]$ by the average log likelihood of the sample, that is

$$L = \frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \pi(y_n = k|z_n, \Theta) \log P(y_n = k|\Theta). \quad (15)$$

Hence the maximization of $\mathbf{E}_\pi[\log P(Y|\Theta)]$ reduces to the maximization of the weighted log likelihood L . As we have defined $P(Y|\Theta)$ by a parametric model previously, we solve the nonlinear optimization problem by the Broyden–Fletcher–Goldfarb–Shanno (BFGS) [46] method. The BFGS method

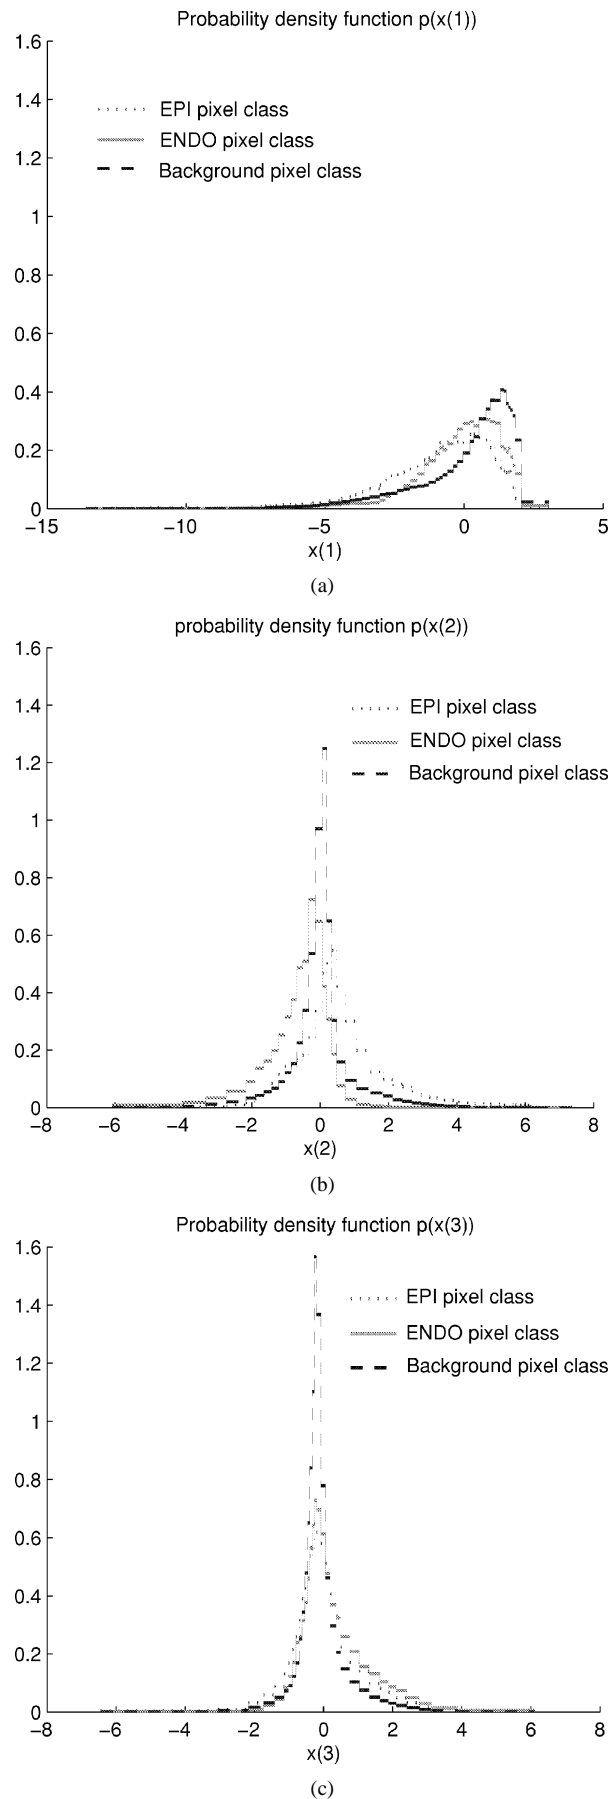


Fig. 6. One-dimensional marginal densities of estimated PixApp probability model.




Fig. 7. Two-dimensional marginal densities of estimated PixApp probability model. (a) EPI. (b) ENDO. (c) Background.




Fig. 8. PixApp probability maps of apical four chamber view. (a)–(c) are for EPI ($p(Z|1)$), ENDO ($p(Z|2)$), and Background ($p(Z|3)$) PixApp maps, respectively. Probabilities are mapped to pixel intensities with histogram equalization.

updates the Hessian matrix by adding a rank-two difference matrix during every iteration and guarantees that the approximated Hessian matrix is positive definite for minimization problems. The major steps include finding the Newton search direction, the line search, and the Hessian update.

Fig. 9 shows the estimated PicPre probability model.

Fig. 10 illustrates the process of pixel class prediction. We obtain a simulated image shown in Fig. 10(a) through ultrasound imaging simulation. Then we compute the distance transform of the visible EPI and ENDO contours, shown as Fig. 10(b) and (c).




Fig. 9. The estimated PicPre probability model.

Fig. 10(d) and (e) shows the intensity maps of the closest on-surface pixels. We apply the estimated PicPre probability model on the distance and intensity maps, and display the PicPre probability maps in Fig. 10(f)–(h).

VI. LEFT VENTRICLE SURFACE MODEL OPTIMIZATION

The goal of surface model optimization is to find an optimal surface model Θ^* that maximizes its posterior probability given all observed images. We use Z_1, Z_2, \dots, Z_N to represent the feature vectors, each corresponding to a pixel. N is the total number of pixels. The optimization problem can be formulated as

$$\max_{\Theta} p(\Theta|Z_1, Z_2, \dots, Z_N) = \max_{\Theta} \frac{p(\Theta)p(Z_1, Z_2, \dots, Z_N|\Theta)}{p(Z_1, Z_2, \dots, Z_N)} \quad (16)$$

Z_1, Z_2, \dots, Z_N themselves are not independent, but strongly statistically related. However, it is reasonable to assume that Z_1, Z_2, \dots, Z_N are statistically independent conditioned upon a given surface model Θ . The interpretation is that given the surface model Θ , the feature vector of a pixel p has no impact on the feature vector of another pixel q statistically. Applying the integrated surface inference theorem, the problem becomes

$$\max_{\Theta} \frac{p(\Theta) \prod_{n=1}^N \sum_{k=1}^K p(Z_n|Y_n = k) P(Y_n = k|\Theta)}{p(Z_1, Z_2, \dots, Z_N)} \quad (17)$$

To simplify the notation, we use $F(\Theta)$ to denote the objective function. Conceptually, this is an unconstrained optimization problem. But considering the implicit prior probability model expressed in the surface model catalog, we still need to modify $F(\Theta)$. Since Θ comes from the linear combinations of the members in the catalog, we can write Θ as

$$\Theta = \sum_{m=1}^M w_m (A_m \Theta_m + T_m)$$

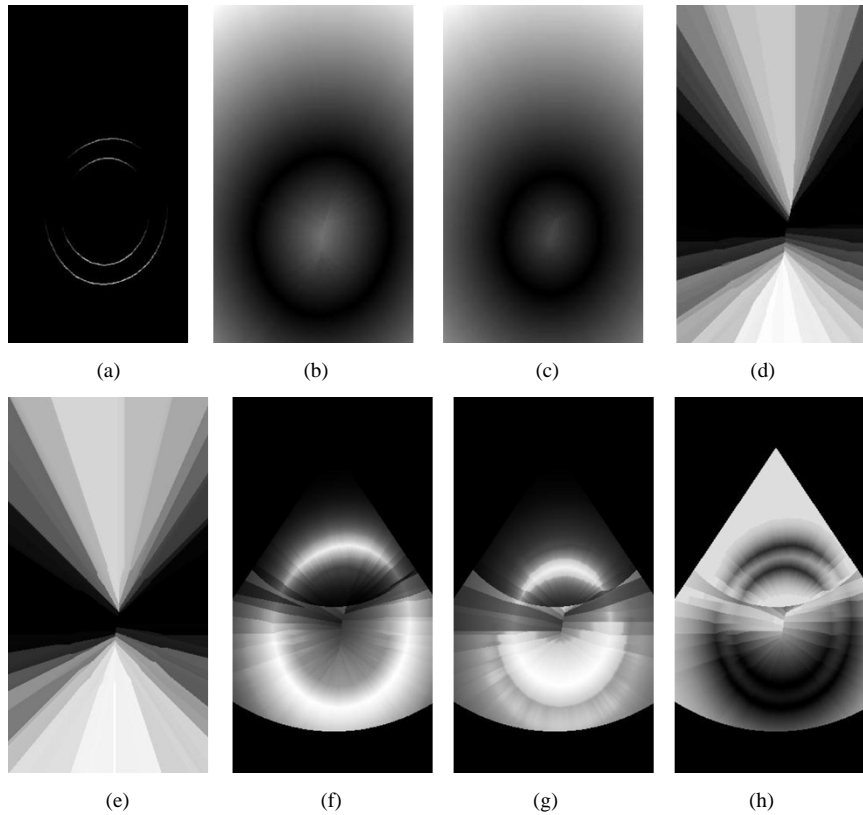


Fig. 10. A simulated image and its intensity, distance and PicPre probability maps. (a) The simulated image. (b), (c) $d(1)$ and $d(2)$ distance maps. (d), (e) $I(1)$ and $I(2)$ intensity maps. (f)–(h) $P(1|\Theta)$, $P(2|\Theta)$, and $P(3|\Theta)$ maps. 1-EPI, 2-ENDO, 3-Background.

where A_m is a transform matrix, T_m is a translation vector, w_m is the weight of Θ_m for Θ , and M is the number of models in the surface model catalog. There might be constraints on w_m , such as all being nonnegative and $\sum_{m=1}^M w_m = 1$. Now the optimization problem is reparameterized into

$$\max_{w_m \geq 0, A_m, T_m, m=1, \dots, M, \sum_{m=1}^M w_m = 1} F \left(\sum_{m=1}^M w_m (A_m \Theta_m + T_m) \right). \quad (18)$$

The constraints $\sum_{m=1}^M w_m = 1$ and $w_m \geq 0$ enforce the combination of the models in the catalog to be convex. Considering the available CPU cycles, our strategy is first to find A_m s and T_m s by an initial alignment of each member in the catalog. Second, A_m s and T_m s are fixed during the optimization of $F(\Theta)$, where only optimal weights are being sought.

In the initial alignment, we compute the transform matrix A_m s and translation vector T_m s for each member in the catalog, such that the aligned surface model maximizes the explicit prior probability given user input points.

After the initial alignment, the A_m s and T_m s are obtained and the catalog is aligned; that is, each member Θ_m in the original catalog is aligned to Θ'_m by $\Theta'_m = A_m \Theta_m + T_m$. Since the problem is still constrained, we reparameterize it by letting

$$w_m = \frac{1 + u_m^2}{M + \sum_{j=1}^M u_j^2} \quad (19)$$

where $u = [u_1, u_2, \dots, u_M] \in \mathbb{R}^M$ is the new optimization variable. The surface model optimization problem can finally be written as

$$\max_{u \in \mathbb{R}^M} F \left(\sum_{m=1}^M \frac{1 + u_m^2}{M + \sum_{j=1}^M u_j^2} \Theta'_m \right). \quad (20)$$

We use the optimization method described in [47, pp. 443–448], which combines both direct search and global methods. It employs the Nelder–Mead simplex algorithm to decide a new search point and determines whether or not to accept the point after randomly perturbing the objective function.

VII. EXPERIMENTAL RESULTS

We evaluate the performance of the integrated approach on 3-D LV surface optimization with a total of 45 *in vivo* clinical studies. There are 16 normal studies and 29 diseased studies. There are six groups among the 45 studies.

- 1) The Normal group consists of healthy studies.
- 2) The CM group consists of cardiomyopathy, idiopathic, and mean of unknown cause studies.
- 3) The IMR group consists of ischemic dilated cardiomyopathy with mitral regurgitation studies.
- 4) The MI group consists of myocardial infarction studies.
- 5) The IO (from University of Iowa) group has one myocardial infarction study.
- 6) The VA (from the VA Hospital) group consists of chronic coronary disease studies.

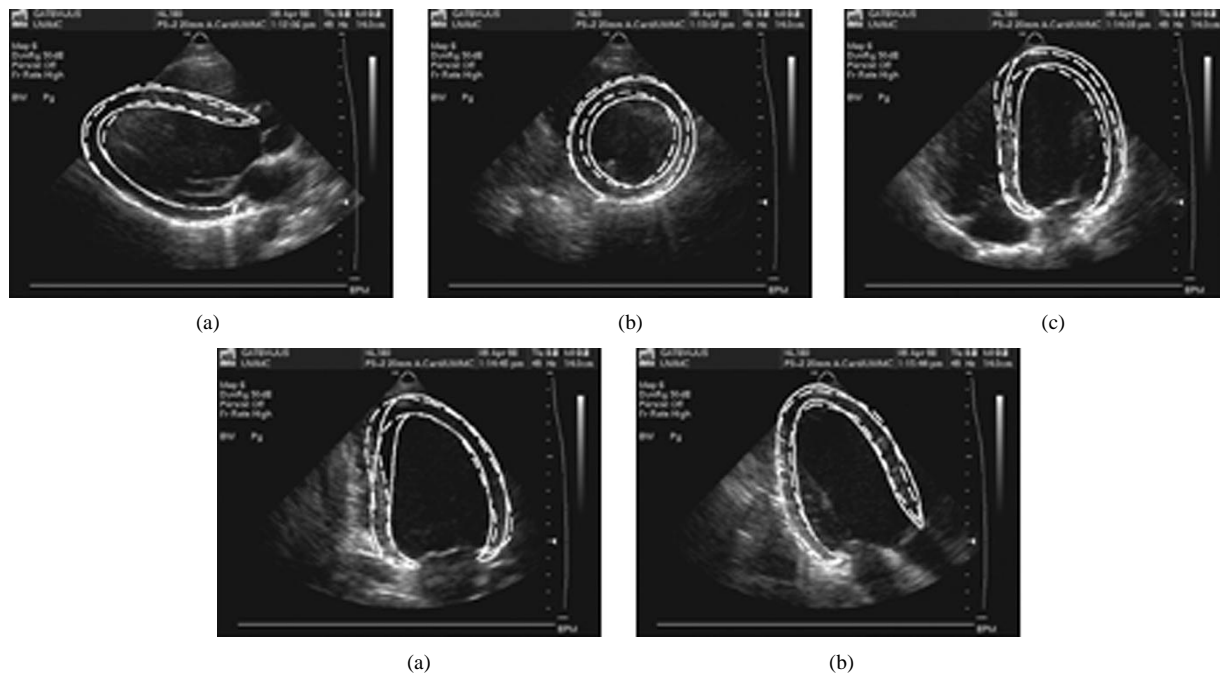


Fig. 11. A normal test study (NL180). Imaging plane intersections of optimized (solid line) and true (dashed line) surface models. (a) Parasternal long-axis view. (b) Short-axis view. (c) Apical four-chamber view. (d) Apical two-chamber view. (e) Apical long-axis view.

The LV surface catalog has 86 LV models from 73 patients: 52 studies in eight diseased groups and 34 normal studies. All the ground-truth LV surface models are built with a labor-intensive and accurate method using more than five views.

All 45 sets of image sequences were acquired from ATL ultrasound machines, except the IO studies, which were acquired from HP ultrasound machines. These images were acquired for other studies by three operators over a period of two years, so that they incorporate some amount of operator and system setting variability. The frame rate was 30 per second. The horizontal and vertical resolutions of the images were, respectively, 0.37 – 0.46 mm and 0.37 – 0.41 mm per pixel. For each of the 45 studies, we selected subsequences of images from four or five different views, including three or four long-axis views and one short-axis view. Each view was further divided into an upper sector and a lower division, divided by an arc passing an inner point of the LV and centered at the transducer location.

From the 45 studies, we selected 20 of good image quality as the training set. We obtained different PixApp and PicPre probability models for each division, all estimated from the training studies. So we had ten pairs of PixApp and PicPre probability models for the five views.

The remaining 25 studies form the test set. Before optimization, a test study is removed from the surface model catalog if it is present. In the initial alignment, we apply rotation and translation only to models in the catalog. During the optimization of weights, we allow a maximum number of 336 objective function evaluations per study. It takes the highly optimized C++/C program on the average about two hours per study on an Intel Pentium III 500-MHz computer to complete the surface model optimization of both ENDO and EPI.

We perform the experiment at *end diastole*. We measure the projection distance between the optimized and the ground-truth surface models. The *projection distance from surface A to sur-*

face B is defined as the mean vertex projection distance from all the vertices of surface *A* to surface *B*. The *projection distance between surface A and B* is the average of the projection distances from *A* to *B* and from *B* to *A*.

Figs. 11 and 12 display the optimization results on a normal group test study and a diseased CM group test study, respectively. Both figures show the original images and the imaging plane intersections with the optimized surface models (solid line) and the ground-truth surface models (dashed line). We achieved surface projection distance errors of 2.0 and 2.2 mm on ENDO and EPI, respectively, on the normal study in Fig. 11. The two surface models agree well at places with strong contrast available, such as the upper and lower center of the view in Fig. 11(a), the lower part of the view in Fig. 11(b), the lower parts of the view in Fig. 11(c), the lower left part of the view in Fig. 11(d), and the left and right parts of Fig. 11(e). However, we observed large error around the area in the top left of the views in Fig. 11(c) and (d). Dropout occurs at these places because the local surfaces there are almost parallel to the incident direction of the ultrasound beam, emitted from the transducer at the origin of the fan-shaped scanning area. In Fig. 12, we had an ENDO surface projection distance of 1.9 mm and an EPI surface projection distance error of 2.5 mm on a diseased CM study. In this study, the error distributed more evenly than in the previous study. The major error occurs around the apex area, which is typically either in the near field where distortions are serious or outside the imaging area. So the apex areas will rely more on the prior shape knowledge and user input points.

Table I gives the surface projection distance error summary statistics on the 25 test studies. The overall distance errors for ENDO and EPI are 2.6 and 3.2 mm, respectively. The best performance is achieved on the normal group test studies. The diseased studies have various errors. Fig. 13 shows the scatter plot of the ENDO/EPI surface projection distance errors of all the

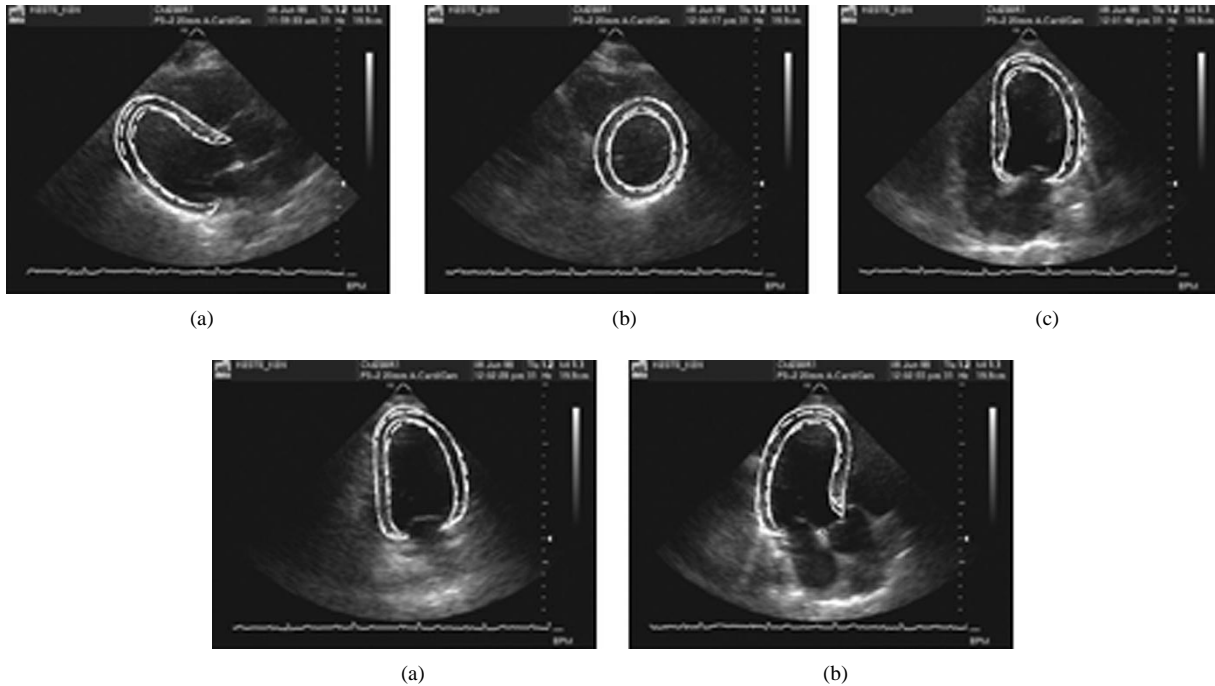


Fig. 12. A diseased test study (CM240R1). Imaging plane intersections of optimized (solid line) and true (dashed line) surface models. (a) Parasternal long-axis view. (b) Short-axis view. (c) Apical four-chamber view. (d) Apical two-chamber view. (e) Apical long-axis view.

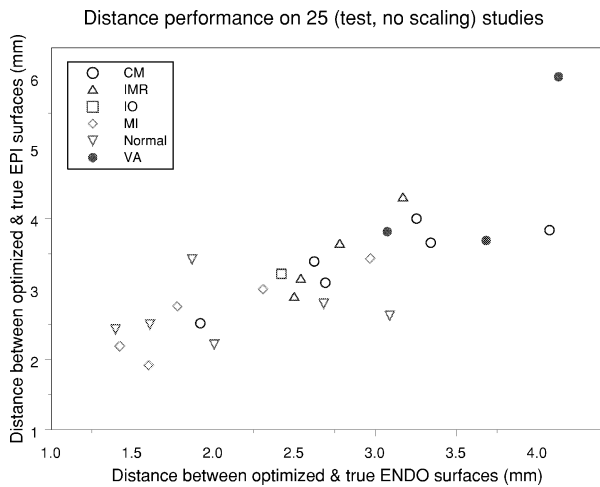


Fig. 13. Scatter plot of surface distance errors on 25 test studies.

test studies. We can see that the normal and MI groups are better than other groups.

We had an IO test study whose images were acquired from an HP ultrasound machine, different from the ATL ultrasound machine used for the 20 training studies. We accomplished a surface projection distance error of 2.4/3.2 mm, smaller than the average error of the 25 test studies. This indicates that our PixApp and PicPre probability models are extendable to capture the underlying imaging process of different types of ultrasound machine.

The optimized surface model for the VA study had larger errors than others because this study had lower than average image quality and the LV had a very abnormal shape. Such cases can be rejected by a lower posterior probability of the optimized surface model given the observed images.

TABLE I
SURFACE DISTANCE ERROR (MM) STATISTICS OF ENDO/EPI OF 25 TEST STUDIES

Group(#Studies)	Mean	Median	Std. Dev.	Min.	Max.
CM(6)	3.0 / 3.4	3.0 / 3.5	0.74 / 0.55	1.9 / 2.5	4.1 / 4.0
IMR(4)	2.7 / 3.5	2.7 / 3.4	0.31 / 0.62	2.5 / 2.9	3.2 / 4.3
IO(1)	2.4 / 3.2	2.4 / 3.2	- / -	2.4 / 3.2	2.4 / 3.2
MI(5)	2.0 / 2.7	1.8 / 2.8	0.63 / 0.61	1.4 / 1.9	3.0 / 3.4
Normal(6)	2.1 / 2.7	1.9 / 2.6	0.65 / 0.42	1.4 / 2.2	3.1 / 3.4
VA(3)	3.6 / 4.5	3.7 / 3.8	0.53 / 1.3	3.1 / 3.7	4.1 / 6.0
Overall(25)	2.6 / 3.2	2.6 / 3.1	0.78 / 0.85	1.4 / 1.9	4.1 / 6.0

TABLE II
STATISTICS OF THE PROJECTION DISTANCE ERROR (MM) BETWEEN USER INPUT SURFACE POINTS AND GROUND-TRUTH SURFACE MODELS. (APEX, AM, PM, ASQ, AND PWM ARE THE NAMES OF THE SURFACE POINTS)

	Mean	Median	Std. Dev.	Min.	Max.
Apex	2.6	2.6	2.1	0.002	8.3
AM	3.0	2.4	2.1	0.3	9.0
PM	2.0	1.5	1.9	0.2	8.8
ASQ	2.3	1.2	2.3	0.4	8.7
PWM	1.8	1.3	1.8	0.1	6.7
Overall	2.3	1.7	2.1	0.002	9.0

We use the projection distance from a user input point to the ground-truth surface model to indicate the human tracing error. Table II shows the user input point tracing error statistics in the 20 training studies. We show five of the seven user input points because these five are supposed to be located on the ENDO surface. They are the apex on ENDO and four ENDO surface points. The mean of the tracing error is 2.3 mm, only 0.3 mm smaller than the mean ENDO surface projection distance error of the 25 test studies. On the other hand, the human tracing error has a much larger standard deviation of 2.1 mm versus 0.78 mm of the 25 test studies. Therefore, even though there is still room

for improvement on the diseased studies, we believe our overall result approaches the achievable performance limit. For clinical use, our approach can be modified to allow the operator to accept or reject the computed boundaries and, if necessary, select additional points to improve the result.

VIII. CONCLUSION

We studied an integrated approach to surface model optimization that incorporates both the low-level image evidence and the high-level prior shape knowledge in finding the best surface model maximizing the posterior probability. This integrated approach falls naturally within the Bayesian framework. It systematically tackles both offline training and online optimization. The offline probability estimation problem is solved by creative use of the EM algorithm.

We applied the integrated approach to 3-D LV surface model optimization from 2-D ultrasound heart images. We showed the qualitative and quantitative results of the optimized LV surface models. Experimental results confirm the feasibility of the integrated approach. The results for normal studies are usually better than the diseased studies because normal LVs are well represented in the catalog. To further improve the performance on the diseased studies, we believe more diseased examples should be put into the surface model catalog. The PixApp and PicPre probability models are jointly estimated. As the probability models do not change over time for a given imaging system, when a dynamic model of the LV is available, the integrated approach still yields an optimal solution and can be applied in tracking 3-D surfaces over time.

Experimental results on 3-D echocardiography have shown that the integrated approach gives a promising direction for further study of surface model optimization. As we are solving an optimization problem whose objective function has a higher dimension and more complexity than the two-stage approach, it demands more CPU cycles both online and offline. Despite the current computational cost, we expect to see a transition from the lower performance two-stage approach to an integrated approach with the ever-increasing CPU speed and decreasing price.

ACKNOWLEDGMENT

The authors would like to thank Dr. J. A. McDonald of Seattle, WA, for suggesting the convex combination approach to creating new surface models. They also appreciate the anonymous reviewers for their comments, which have improved the overall quality of this paper.

REFERENCES

- [1] M. E. Legget, D. F. Leotta, E. L. Bolson, J. A. McDonald, R. W. Martin, X.-N. Li, C. M. Otto, and F. H. Sheehan, "System for quantitative three-dimensional echocardiography of the left ventricle based on a magnetic-field position and orientation sensing system," *IEEE Trans. Bio. Eng.*, vol. 45, pp. 494–504, Apr. 1998.
- [2] M. Song and R. M. Haralick, "Optimal grid quantization," in *Proc. Int. Conf. Pattern Recognition*, vol. III, Quebec City, Canada, Aug. 2002, pp. 444–447.
- [3] A. F. Frangi, W. J. Niessen, and M. A. Viergever, "Three-dimensional modeling for functional analysis of cardiac images: A review," *IEEE Trans. Med. Imag.*, vol. 20, pp. 2–25, Jan. 2001.

- [4] H. Hoppe, T. DeRose, T. Duchamp, M. Halstead, H. Jin, J. McDonald, J. Schweitzer, and W. Stuetzle, "Piecewise smooth surface reconstruction," in *Proc. SIGGRAPH '94*, July 1994, pp. 295–302.
- [5] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle, "Surface reconstruction from unorganized points," in *Comput. Graph. (Proc. SIGGRAPH '92)*, vol. 26, July 1992, pp. 71–78.
- [6] —, "Mesh optimization," in *Proc. SIGGRAPH '93*, Aug. 1993, pp. 19–26.
- [7] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active contour models," *Int. J. Comput. Vision*, vol. 1, pp. 321–331, Jan. 1987.
- [8] V. Chalana, D. T. Linker, D. R. Haynor, and Y. Kim, "A multiple active contour model for cardiac boundary detection on echocardiographic sequences," *IEEE Trans. Med. Imag.*, vol. 15, pp. 290–298, June 1996.
- [9] D. C. Wilson, E. A. Geiser, and J.-H. Li, "Feature extraction in two-dimensional short-axis echocardiographic images," *J. Math. Imag. Vision*, vol. 3, pp. 285–298, 1993.
- [10] D. C. Wilson, E. A. Geiser, D. A. Conetta, J. M. Murphy, and D. Wang, "An automated algorithm for analysis of 2-D echocardiographic short-axis images: A brief overview," in *Proc. Mathematical Methods in Biomedical Image Analysis*, 1996, pp. 222–231.
- [11] E. A. Geiser, D. C. Wilson, D. X. Wang, D. A. Conetta, J. D. Murphy, and A. D. Hutson, "Autonomous epicardial and endocardial boundary detection in echocardiographic short-axis images," *J. Amer. Soc. Echocardiogr.*, vol. 11, no. 4, pp. 338–348, Apr. 1998.
- [12] J. F. Brinkley, "Knowledge-driven ultrasonic three-dimensional organ modeling," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-7, pp. 431–441, July 1985.
- [13] G. D. Stetten and S. M. Pizer, "Automated identification and measurement of objects via populations of medial primitives, with application to real time 3D echocardiography," in *Information Processing in Medical Imaging '99*, Lecture Notes in Computer Science, 1999, pp. 84–97.
- [14] —, "Medial-node models to identify and measure objects in real-time 3-D echocardiography," *IEEE Trans. Med. Imag.*, vol. 18, pp. 1025–1034, Oct. 1999.
- [15] A. Pentland, "Recognition by parts," in *Proc. IEEE Int. Conf. Computer Vision*, 1987, pp. 612–620.
- [16] F. Solina and R. Bajcsy, "Recovery of parametric models from range images: The case for superquadrics with global deformations," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 12, pp. 131–147, Feb. 1990.
- [17] E. Bardinet, L. D. Cohen, and N. Ayache, "Analyzing the deformation of the left ventricle of the heart with a parametric deformable model," INRIA, Paris, France, Tech. Rep. 2797, 1996.
- [18] J. Park, D. Metaxas, and L. Axel, "Analysis of left ventricular wall motion based on volumetric deformable models and MRI-SPAMM," *Med. Image Anal.*, vol. 1, no. 1, pp. 53–71, 1996.
- [19] J. F. Brinkley, "A flexible, generic model for anatomic shape: Application to interactive two-dimensional medical image segmentation and matching," *Comput. Meth. Programs Biomed.*, vol. 26, pp. 121–142, 1993.
- [20] R. B. Altman and J. F. Brinkley, "Probabilistic constraint satisfaction with structural models: Application to organ modeling by radial contours," in *17th Annu. Symp. Computer Applications in Medical Care*, Washington, DC, 1993, pp. 492–496.
- [21] T. F. Cootes, A. Hill, C. J. Taylor, and J. Haslam, "The use of active shape models for locating structures in medical images," *Image Vision Comput.*, vol. 12, no. 6, pp. 355–366, July 1994.
- [22] T. F. Cootes, G. J. Edwards, and C. J. Taylor, "Active appearance models," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 23, pp. 681–685, June 2001.
- [23] R. H. Davies, T. F. Cootes, and C. J. Taylor, "A minimum description length approach to statistical shape modeling," in *Information Processing Medical Imaging*, vol. 2082, Lecture Notes in Computer Science, 2001, pp. 50–63.
- [24] A. Blake and M. Isard, *Active Contours: The Application of Techniques From Graphics, Vision, Control Theory and Statistics to Visual Tracking of Shapes in Motion*. Berlin, Germany: Springer, 1998.
- [25] A. Blake, B. Bascle, M. Isard, and J. MacCormick, "Statistical models of visual shape and motion," *Phil. Trans. Roy. Soc. London A*, vol. 356, pp. 1283–1302, 1998.
- [26] G. Jacob, J. A. Noble, and A. Blake, "Robust contour tracking in echocardiographic sequences," in *Proc. 6th Int. Conf. Computer Vision*, Bombay, India, Jan. 1998, pp. 408–413.
- [27] A. Cauce and C. J. Taylor, "3D point distribution models of the cortical sulci," in *Proc. 6th Int. Conf. Computer Vision*, New Delhi, India, 1998, pp. 402–406.

- [28] C. Lorenz and N. Krahnstöver, "Generation of point-based 3D statistical shape models for anatomical objects," *Comput. Vision Image Understand.*, vol. 77, no. 2, pp. 175–191, 2000.
- [29] Y. H. Tseng, J. N. Hwang, and F. H. Sheehan, "3-D heart contour delineation and motion estimation of ultrasound images using continuous distance transform neural networks," *J. VLSI Signal Process.*, vol. 18, no. 3, pp. 207–218, 1998.
- [30] I. L. Dryden and K. V. Mardia, *Statistical Shape Analysis*, ser. Wiley Series in Probability and Statistics. New York: Wiley, 1998.
- [31] S. R. Lele and J. T. Richtsmeier, *An Invariant Approach to Statistical Shape Analysis*. New York: Chapman & Hall/CRC, 2001.
- [32] J. Canny, "A computational approach to edge detection," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-8, pp. 679–698, Nov. 1986.
- [33] R. M. Haralick and L. G. Shapiro, *Computer and Robot Vision*. Reading, MA: Addison-Wesley, 1992, vol. I.
- [34] G. Coppini, R. Poli, and G. Valli, "Recovery of the 3-D shape of the left ventricle from echocardiographic images," *IEEE Trans. Med. Imag.*, vol. 14, pp. 301–317, June 1995.
- [35] A. Chakraborty, L. H. Staib, and J. S. Duncan, "Deformable boundary finding in medical images by integrating gradient and region information," *IEEE Trans. Med. Imag.*, vol. 15, pp. 859–870, Dec. 1996.
- [36] J. Montagnat, H. Delingette, and G. Malandain, "Cylindrical echocardiographic image segmentation based on 3D deformable models," in *Medical Image Computing Computer-Assisted Intervention*, vol. 1679, Lecture Notes in Computer Science, 1999, pp. 168–175.
- [37] M. Mignotte and J. Meunier, "Deformable template and distribution mixture-based data modeling for the endocardial contour tracking in an echographic sequence," in *Proc. Computer Vision Pattern Recognition*, 1999, pp. 225–230.
- [38] R. T. Cox, "Probability, frequency and reasonable expectation," *Amer. J. Phys.*, vol. 14, pp. 1–13, 1964.
- [39] E. T. Jaynes, "Bayesian methods: General background," in *Maximum Entropy and Bayesian Methods in Statistics*, J. H. Justice, Ed. Cambridge, U.K.: Cambridge Univ. Press, 1986, pp. 1–25.
- [40] M. Song, "Ultrasound imaging simulation and echocardiographic image synthesis," Master's thesis, Dept. of Electrical Engineering, Univ. of Washington, Seattle, 1999.
- [41] M. Song, R. M. Haralick, and F. H. Sheehan, "Ultrasound imaging simulation and echocardiographic image synthesis," in *Proc. Int. Conf. Image Processing*, vol. III, Vancouver, Canada, Sept. 2000, pp. 420–423.
- [42] M. J. Seaidoun, "A fast exact Euclidean distance transform with application to computer vision and digital image processing," Ph.D. dissertation, Northeastern Univ., 1993.
- [43] T. Hebert and R. Leahy, "A generalized EM algorithm for 3-D Bayesian reconstruction from poisson data using Gibbs priors," *IEEE Trans. Med. Imag.*, vol. 8, pp. 194–202, 1989.
- [44] P. J. Green, "Bayesian reconstruction from emission tomography data using a modified EM algorithm," *IEEE Trans. Med. Imag.*, vol. 9, pp. 84–93, 1990.
- [45] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. Roy. Statist. Soc. B*, vol. 39, no. 1, pp. 1–38, 1977.
- [46] J. Nocedal and S. J. Wright, *Numerical Optimization*, ser. Springer Series in Operations Research. New York: Springer-Verlag, 1999.
- [47] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in Fortran 77: The Art of Scientific Computing*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 1996, vol. 1.