A New Deformable Model for Analysis of X-Ray CT Images in Preclinical Studies of Mice for Polycystic Kidney Disease

S. S. Gleason*, *Member, IEEE*, H. Sari-Sarraf, *Member, IEEE*, M. A. Abidi, *Member, IEEE*, O. Karakashian, and F. Morandi

***Abstract*—This paper describes the application of a new probabilistic shape and appearance model (PSAM) algorithm to the task of detecting polycystic kidney disease (PKD) in X-ray computed tomography images of laboratory mice. The genetically engineered PKD mouse is a valuable animal model that can be used to develop new treatments for kidney-related problems in humans. PSAM is a statistical-based deformable model that improves upon existing point distribution models for boundary-based object segmentation. This new deformable model algorithm finds the optimal boundary position using an objective function that has several unique characteristics. Most importantly, the objective function includes both global shape and local gray-level characteristics, so optimization occurs with respect to both pieces of information simultaneously. PSAM is employed to segment the mouse kidneys and then texture measurements are applied within kidney boundaries to detect PKD. The challenges associated with the segmentation nonrigid organs along with the availability of a priori information led to the choice of a trainable, deformable model for this application. In 103 kidney images that were analyzed as part of a preclinical animal study, the mouse kidneys and spine were segmented with an average error of 2.4 pixels per boundary point. In all 103 cases, the kidneys were successfully segmented at a level where PKD could be detected using mean-of-local-variance texture measurements within the located boundary.**

***Index Terms*—Deformable models, screening, segmentation, statistical shape models.**

I. Introduction: Medical Imaging and Deformable Models

THE Oak Ridge National Laboratory (ORNL), Oak Ridge, TN, has a Mammalian Genetics Research Facility that houses more than 70 000 mice representing about 400 mutant lines. Mutagenesis experiments are performed on the mice, and it is important to then determine the physical manifestations (phenotypes) of these induced mutations. These phenotypes are often difficult to find, especially when only a few researchers are available to screen a large number of mice. A large percentage of these phenotypes are expressed as internal abnormalities that cannot be seen without sacrificing the animal. Examples of these phenotypes include skeletal deformities (e.g., scoliosis) and deformed or diseased organs [e.g., polycystic kidney disease (PKD)]. In support of this phenotyping research, an X-ray micro-computed tomography (CT) system (the MicroCAT) was developed that can acquire CT data sets with a resolution as high as 50 microns [1]. MicroCAT system generates large volumetric data sets of the mouse anatomy that must be analyzed to determine if a particular phenotype is present in a given subject. To perform phenotypic screening on large numbers of mice, an automated approach is needed to identify potential anatomic mutations within the animals.

This paper describes the application of a new statistical-based deformable model algorithm to the segmentation of kidneys in X-ray CT images of laboratory mice. This segmentation algorithm has been developed as the crucial first step in a process to automatically screen mice for genetically-induced PKD. Once the segmentation is complete, mean-of-local-variance (MOLV) texture measurements are applied within detected kidney boundaries to discern the presence of PKD. This new algorithm is called probabilistic shape and appearance model (PSAM) and has its roots in the active shape model (ASM) algorithm developed by Cootes, *et al.* [2]. ASM was chosen as an appropriate starting point for the development of a segmentation algorithm for this application because of its applicability to segmentation problems with the following characteristics:

- shape is a primary characteristic that describes the object;
- images contain complex backgrounds;
- objects have faint, obscured, or partially missing object boundaries;
- there exists available a priori information on object appearance.

ASM incorporates *a priori* information extracted from a training set to build a gray-level model (GLM) and a global shape model (GSM). These models are used during an iterative contour deformation process that adjusts the position and shape of the contour to match the boundary of the object within the image. Although ASM is an excellent starting point for the motivating application, it has a shortcoming in a key area: global shape

Manuscript received November 15, 2001; revised August 1, 2002. This work was supported by the Laboratory Director's Research and Development Program at the Oak Ridge National Laboratory managed by UT-Battelle, LLC, for the U.S. Department of Energy. *Asterisk indicates corresponding author.*

*S. S. Gleason is with the Engineering Science and Technology Division, Oak Ridge National Laboratory, Oak Ridge, TN 37831-2008 USA (e-mail: gleasonss@ornl.gov).

H. Sari-Sarraf is with the Electrical Engineering Department, Texas Tech University, Lubbock, TX 79409 USA.

M. A. Abidi is with the Electrical Engineering Department, University of Tennessee, Knoxville, TN 37996 USA.

O. Karakashian is with the Mathematics Department, University of Tennessee, Knoxville, TN 37996 USA.

F. Morandi is with the Department of Small Animal Clinical Sciences, University of Tennessee, Knoxville, TN 37996 USA.

Digital Object Identifier 10.1109/TMI.2002.806278

and gray-level information are treated independently during optimization of the boundary position. This shortcoming limits ASM's robustness and accuracy in some applications, including the kidney screening application presented here. A set of kidney cross-sectional images from 12 different mice is shown in Fig. 1. These images were acquired using the MicroCAT X-ray CT system at a 200 μm pixel resolution. The mice were given intraperitoneal injections of 200 μl of iodinated contrast medium (Conray 400) 23–25 min prior to CT scanning. Note the variability in shape, density, pose, and contrast of the kidneys (relative to their backgrounds) among the 12 images. Also note the varying complexities in the background structure. In particular, note the faint boundaries around many of the kidneys (e.g., second row, middle image).

Before describing the details of the new PSAM algorithm and its application to this problem, it is appropriate to highlight recent developments in statistical-based deformable models that are variations on ASM. Wang, *et al.* [3] have developed a probabilistic based optimization scheme that uses Cootes' point distribution model and integrates Canny-edge information into the maximum *a posteriori* (MAP) objective function as the underlying image attraction force. Kervrann, *et al.* [4] have developed similar probabilistic techniques, but include Markov modeling on the local scale to promote boundary smoothness. Both of these ASM adaptations rely on edge information in the target image as the external attraction force, rather than the gray-level gradients proposed by Cootes. Duta, *et al.* [5] have also refined the ASM technique in terms of the image attraction force as well as the optimization approach to fit the boundary to the underlying image data. Consideration of boundary-point outliers is an important consideration in their work. Gleason, *et al.* [6] have improved upon the original ASM by including more comprehensive gray-level information from the image (raw intensity plus profile gradients) and have added constraints to the shape model to improve convergence. Even with the considerable research that has been performed, none of the resulting approaches address the key ASM shortcoming outlined previously of independent optimization with respect to gray-level and shape information.

Another algorithm that is an extension of the idea behind ASM is the active appearance model (AAM) [7]. In ASM the boundary of the object is modeled using a set of landmark points (LPs) that define the exterior boundary of the structure, while AAM models the entire object appearance as a collection of gray-level values. It turns out that AAM is not applicable to the problem of PKD screening, because PKD causes significant changes in the internal kidney appearance. As a result, it makes sense to only model the part of the kidney where the relative intensity characteristics are more consistent, that is, the boundary. For this reason, a new PDM technique was developed that takes advantage of the modeling approach of ASM, but improves upon that technique with a new boundary optimization approach.

II. Shape and Gray-Level Objective Function

This section outlines the theoretical development of the PSAM segmentation algorithm. First, we will define the

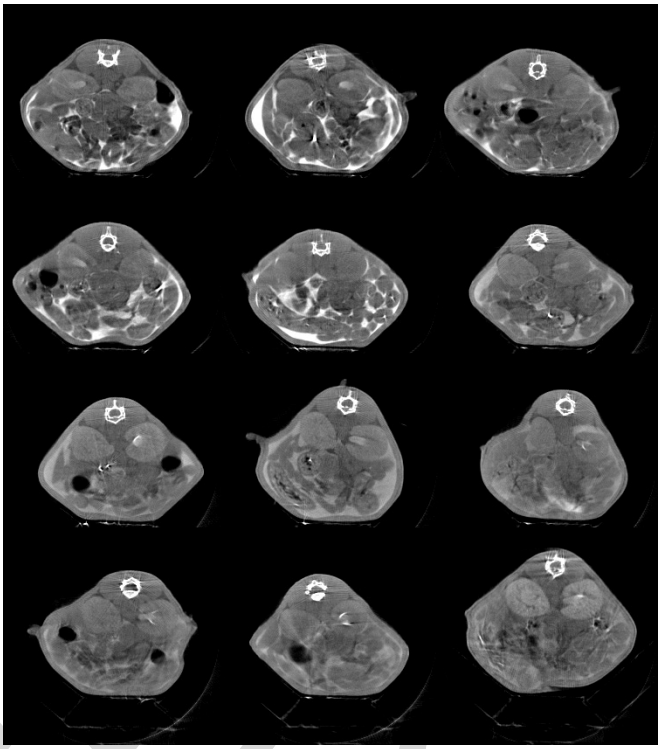

Fig. 1. X-ray CT abdominal cross-sections of 12 different mice. The white structure at the top of each image is the spinal column. The two elliptical shaped structures to the lower-left and right of the spinal column are the kidneys. The black spots in the abdomen are air pockets and the distributed white abdominal structure (especially in the first five images) is excess contrast agent.

boundary we are searching for as an (x, y) coordinate matrix, $\boldsymbol{S_v}$, of size $N \times 2$, where N is the number of LPs needed to represent the boundary. Each row of $\boldsymbol{S_v}$ corresponds to the (x, y) coordinate of a boundary or LP as follows:

$$\boldsymbol{S_v} = \begin{bmatrix} x_1 & y_1 \\ x_2 & y_2 \\ \dots & \dots \\ x_N & y_N \end{bmatrix} \tag{1}$$

where the N LPs are defined by the coordinate pairs

$$(x_j, y_j), j = 1, \dots N. \tag{2}$$

The subscript $\boldsymbol{v}$ is used to foreshadow the relationship of the boundary with a combined shape-pose parameter vector, $\boldsymbol{v}$, to be described later. Defining the boundary this way accommodates some flexibility in how the boundary is represented. The points in $\boldsymbol{S_v}$ are typically LPs that lie directly on the boundary (as used in ASM), but in addition, other (x, y) coordinate locations could be included that, for example, lie in image regions that are on the interior or exterior of the object boundary.

Next, we will define a feature vector matrix, $\boldsymbol{G}$. In this formulation, the feature vectors (rows of $\boldsymbol{G}$) contain features extracted from the neighborhood of each LP on the boundary. If the length of each feature vector is m, then the size of $\boldsymbol{G}$ will be $N \times m$. These feature vectors may contain any image information that is relevant to the given application, such as gray-level intensities in the neighborhood of each pixel, local texture measurements,

or gray-level gradient information (as used in ASM and PSAM). We can write $\boldsymbol{G}$ as

$$\boldsymbol{G} = \begin{bmatrix} \boldsymbol{g}_1 \\ \boldsymbol{g}_2 \\ \cdots \\ \boldsymbol{g}_N \end{bmatrix} \tag{3}$$

where

$$\boldsymbol{g}_j = h(f(x_j, y_j)), j = 1, \ldots, N. \tag{4}$$

In (4), f is the image, and h is a neighborhood operator that samples the image in the neighborhood of (x_j, y_j) and organizes these samples into the feature vector, $\boldsymbol{g}_j$. Now that we have defined our parameters, we can formulate the probabilistic objective function.

In the PSAM boundary-finding application, the goal is to maximize the a posteriori probability of the boundary $\boldsymbol{S_v}$, given the measured image features in $\boldsymbol{G}$. Using the compound version of Bayes rule [8], we can write the *a posteriori* probability expression for the boundary, $\boldsymbol{S_v}$, given a collection of feature vectors, $\boldsymbol{G}$, as

$$P(\boldsymbol{S_v}|\boldsymbol{G}) = \frac{p(\boldsymbol{G}|\boldsymbol{S_v})P(\boldsymbol{S_v})}{p(\boldsymbol{G})} \tag{5}$$

where $P(\boldsymbol{S_v})$ is the prior probability of a boundary instance, $\boldsymbol{S_v}$; $p(\boldsymbol{G}|\boldsymbol{S_v})$ is the conditional density of $\boldsymbol{G}$ given the boundary instance, $\boldsymbol{S_v}$; and $p(\boldsymbol{G})$ is the prior probability density for $\boldsymbol{G}$. The goal is then to optimize (5) by searching over all possible values of $\boldsymbol{S_v}$ to find the one, $\boldsymbol{S_v^*}$, that corresponds to the MAP value as given by

$$P(\boldsymbol{S_v^*}|\boldsymbol{G}) = \arg\max_{\boldsymbol{S_v}} \left\{ \frac{p(\boldsymbol{G}|\boldsymbol{S_v})P(\boldsymbol{S_v})}{p(\boldsymbol{G})} \right\}. \tag{6}$$

Finding $\boldsymbol{S_v^*}$ can be further simplified by writing the objective function [we will now call it $J(\boldsymbol{S_v})$] can be written as

$$J(\boldsymbol{S_v}) = \ln p(\boldsymbol{G}|\boldsymbol{S_v}) + \ln P(\boldsymbol{S_v}) \tag{7}$$

where we have dropped the term $-\ln p(\boldsymbol{G})$ since it is independent of $\boldsymbol{S_v}$ and is therefore a constant.

Depending on the application at hand, the term $\ln p(\boldsymbol{G}|\boldsymbol{S_v})$ may be difficult to calculate. If independence is assumed between the feature vectors in $\boldsymbol{G}$, and if it is also assumed that each feature vector, $\boldsymbol{g}_j$, is dependent only on it corresponding location, (x_j, y_j) (i.e., jth row of $\boldsymbol{S_v}$), then the conditional density of $\boldsymbol{G}$ given $\boldsymbol{S_v}$ can be rewritten as [8]

$$p(\boldsymbol{G}|\boldsymbol{S_v}) = \prod_{j=1}^{N} p(\boldsymbol{g}_j|S_{v_{j,1}}, S_{v_{j,2}}) = \prod_{j=1}^{N} p(\boldsymbol{g}_j|x_j, y_j). \tag{8}$$

Plugging this expression back into (7), we can write

$$\begin{aligned} J(\boldsymbol{S_v}) &= \sum_{j=1}^{N} \ln p(\boldsymbol{g}_j|S_{v_{j,1}}, S_{v_{j,2}}) + \ln P(\boldsymbol{S_v}) \\ &= \sum_{j=1}^{N} \ln p(\boldsymbol{g}_j|x_j, y_j) + \ln P(\boldsymbol{S_v}). \end{aligned} \tag{9}$$

It is useful to note that the first term is the "data-driven" term of the objective function in that it depends on image characteristics (external energy term), while the second term is "model-driven" in that it depends on prior distributions of boundary shape and location (internal energy term).

Note how this formulation can accommodate several important goals. First, this objective function allows simultaneous optimization with respect to image-derived gray-scale information (first term) and shape information (second term). Second, because this objective function is based on a probabilistic framework, we can interpret its value as a measure of how well the final boundary fits the distribution approximated by those contained in the training set. This measure can be broken down into two pieces: 1) the first term measures how well the final gray-level information matches that which was extracted from the training data, and 2) the second term measures how well the overall shape and location of the boundary matches that which was extracted from the training set.

Optimizing $J(\boldsymbol{S_v})$ over all possible boundary vectors, $\boldsymbol{S_v}$, can be a daunting task for several reasons. First, depending on the number of LPs used to represent the boundary, $\boldsymbol{S_v}$ could be a very large matrix, and finding the maximum of J with respect to each of the $2N$ elements of $\boldsymbol{S_v}$ can be computationally demanding. Applying principal component analysis (PCA) and then using only the significant modes of variation reduces the dimensionality of $\boldsymbol{S_v}$ to resolve this problem. The boundary $\boldsymbol{S_v}$ can be approximated in the PCA subspace as a vector, $\boldsymbol{v}$, with fewer dimensions. As detailed in [9], $\boldsymbol{v}$ is constructed to be a combination of the boundary coordinates in PCA subspace, $\boldsymbol{b}$, and the pose of the boundary, $\boldsymbol{z}$, relative to the mean shape calculated during the shape training process (see next section) as follows:

$$\boldsymbol{v} = \begin{bmatrix} \boldsymbol{b} \\ \boldsymbol{z} \end{bmatrix} \tag{10}$$

where $\boldsymbol{z} = [\, s \quad \theta \quad T_x \quad T_y \,]^T$. Here, s is scale, θ is rotation, and T_x, T_y are the x- and y-translations required to align the boundary with the mean shape. Also, we define the length of the PCA shape vector, $\boldsymbol{b}$, to be t_s, the length of $\boldsymbol{z}$ is 4, and the overall length of $\boldsymbol{v}$ is then $t = t_s + 4$.

Hence, optimizing the objective function in the PCA subspace with respect to the more compact vector, $\boldsymbol{v}$, is a simpler task. If we substitute the new boundary representation, $\boldsymbol{v}$, into (7), the objective function then takes the form

$$J(\boldsymbol{v}) = \ln p(\boldsymbol{G}|\boldsymbol{v}) + \ln P(\boldsymbol{v}) \tag{11}$$

where

$$\begin{aligned} \ln p(\boldsymbol{G}|\boldsymbol{v}) &= \sum_{J=1}^{N} \ln p(\boldsymbol{g}_j|q(\boldsymbol{v})_j, r(\boldsymbol{v})_j) \\ &= \sum_{j=1}^{N} \ln p(\boldsymbol{g}_j|x_j, y_j) \end{aligned} \tag{12}$$

and $q(\boldsymbol{v})$, $r(\boldsymbol{v})$ are functions that map $\boldsymbol{v}$ from the PCA subspace back into the x and y image-coordinates, respectively.

To use this objective function in practice, we must know all of the individual conditional probability densities, $p(\boldsymbol{g}_j|x_j, y_j)$,

$j = 1,\ldots,N$. In the PSAM approach, these conditional densities can be straightforwardly measured from the training data as described later. Finally, the prior probability of the PCA boundary vector, $P(\boldsymbol{v})$, must also be known. Once again, this can be estimated from the training data.

A. Shape Model Training

Although it was convenient to represent the boundary as a matrix, $\boldsymbol{S_v}$, during the compound Bayesian description of the objective function, it is simpler for implementation to represent the entire collection of (x, y) coordinates as a vector, $\boldsymbol{p}$. Let the manually selected boundary for the ith image in an M-image training set be represented by a collection of LPs, $\boldsymbol{p}_{train_i}$ as

$$\boldsymbol{p}_{train_i} = [\, x_{i_1} \quad x_{i_2} \quad \ldots \quad x_{i_N}\ y_{i_1} \quad y_{i_2} \quad \ldots \quad y_{i_N}\,], \qquad i = 1,\ldots,M \quad (13)$$

where the N LPs that make up the boundary for the ith sample are defined by the coordinate pairs

$$(x_{i_j}, y_{i_j}), j = 1,\ldots,N. \quad (14)$$

These training boundaries are manually delineated by a qualified person that understands the characteristics of the object within the scene that must be segmented and recognized.

Prior to building the shape model, the alignment of the sets of LPs to a common coordinate frame is done via Procrustes analysis [10] to form the aligned sets of LPs

$$\hat{\boldsymbol{p}}_{train_i} = [\, \hat{x}_{i_1} \quad \hat{x}_{i_2} \quad \ldots \quad \hat{x}_{i_N}\ \hat{y}_{i_1} \quad \hat{y}_{i_2} \quad \ldots \quad \hat{y}_{i_N}\,] \qquad i = 1,\ldots,M. \quad (15)$$

Using this training data set of aligned LPs, $\hat{\boldsymbol{p}}_{train_i}$, we can straightforwardly create a GSM based on $\boldsymbol{v}$ by applying PCA as described in [9].

B. Gray-Level Model Training

Recall that $\boldsymbol{g}_j$ is a vector of image gray-level values in the neighborhood of the jth LP. For this algorithm we assume that the values of $\boldsymbol{g}_j$ lie along a linear profile that passes through the jth LP and is normal to the current boundary estimate. We can write the raw pixel intensity elements of the normal profile as

$$g_{j_k} = f\left(x_j - \left(\frac{(N_g-1)}{2} - k + 1\right)\cos\alpha_j,\ y_j - \left(\frac{(N_g-1)}{2} - k + 1\right)\sin\alpha_j\right) \quad k = 1,\ldots,N_g \text{ and } j = 1,\ldots,N \quad (16)$$

where N_g is the number of gray-level samples in each profile, and α_j is the angle of the profile through the jth LP normal to the boundary. For this application, the elements of the vector $\boldsymbol{g}_j$ are not the raw image intensity values, but the gradient along the profile to make the values independent of global intensity. Once we have extracted the gray level gradient profile vectors for each LP on the boundary for all M images in the training set, we can (similar to the GSM) create a GLM using PCA as described in detail in [9].

C. Objective Function Parameterization and Optimization

Once the formulation of the GSM and GLM is complete, all of the required information is available to parameterize and optimize the objective function in (11). We can rewrite (11) as functions of J_1 and J_2 as follows:

$$J(v) = J_1(v) + J_2(v) \quad (17)$$

where $J_1(v) = \ln P(\boldsymbol{v})$ and $J_2(v) = \ln p(\boldsymbol{G}|\boldsymbol{v})$. For parameterization purposes, the distributions of both the shape vectors ($\hat{\boldsymbol{p}}_{train_i}$) and the gray-level vectors for each LP ($\boldsymbol{g}_j$) are assumed to be Gaussian as in [6] and [9]. Gradient descent was chosen to optimize the objective function. The gradient can be written as

$$\nabla J(\boldsymbol{v}) = \nabla J_1(\boldsymbol{v}) + J_2(\boldsymbol{v}). \quad (18)$$

Because we assume Gaussian distributions for both the GSM and GLM, the gradient of each term can be calculated. The gradient of the shape term is given by

$$\frac{\partial}{\partial v_n} J_1(\boldsymbol{v}) = \frac{2(v_n - \bar{v}_n)}{\sigma_n^2}, n = 1,\ldots,t \quad (19)$$

where σ_n^2 is the variance of v_n as measured from the training set.

The gray-level model term is more complex. The gradient of J_2 is given by

$$\frac{\partial}{\partial v_n} J_2(\boldsymbol{v}) = 2\sum_{j=1}^{N} \dot{\boldsymbol{g}}_j(\boldsymbol{v})^T \boldsymbol{K}_{g_j} \frac{\partial}{\partial v_n}\dot{\boldsymbol{g}}_j(\boldsymbol{v}). \quad (20)$$

Note that $\dot{\boldsymbol{g}}_j$, the profile with the mean subtracted (see (21), is being written as $\dot{\boldsymbol{g}}_j(\boldsymbol{v})$ to indicate that it is a function of $\boldsymbol{v}$, the PCA-based boundary. The partial derivative term in (20) is given by

$$\frac{\partial}{\partial v_n}\dot{\boldsymbol{g}}_j(\boldsymbol{v}) = \frac{\partial}{\partial v_n}(\boldsymbol{g}_j(\boldsymbol{v}) - \bar{\boldsymbol{g}}_j) = \frac{\partial}{\partial v_n}\boldsymbol{g}_j(\boldsymbol{v}) \quad (21)$$

where $\bar{\boldsymbol{g}}_j$ is the mean gray-level profile through the jth LP. The partial derivative of $\boldsymbol{g}_j$ depends, of course, on how the gray-level profiles are defined during training, as well as on the operator, $h(\boldsymbol{g}_j)$. The profile samples here are calculated as edge profiles based on a first-forward difference gradient

$$g_{j_k}(\boldsymbol{v}) = \boldsymbol{I}(Q_{j_k}(\boldsymbol{v}), R_{j_k}(\boldsymbol{v})) - \boldsymbol{I}(Q_{j_{k+1}}(\boldsymbol{v}), R_{j_{k+1}}(\boldsymbol{v})), \qquad k = 1,\ldots,N_g \quad (22)$$

Here, $\boldsymbol{Q}_j$ and $\boldsymbol{R}_j$ are vectors that contain the values of the (x, y) coordinate locations of the samples along the jth normal profile and are defined as [see (16)]

$$Q_{j_k} = x_j - \left(\frac{(N_g-1)}{2} - k + 1\right)\cos\alpha_j \qquad k = 1,\ldots,N_g \quad (23)$$

$$R_{j_k} = y_j - \left(\frac{(N_g-1)}{2} - k + 1\right)\sin\alpha_j \qquad k = 1,\ldots,N_g. \quad (24)$$

where α_j is the angle of the jth profile. Also note that $\boldsymbol{I}$ is the image under test and that (q_j, r_j) is the pose-corrected coordinate pair indicating the location of the jth LP within the image.

The profile is also normalized by the sum of the absolute values of the profile elements, $\sum_{m=1}^{N} |g_{j_m}|$. The expression for the profile in this case is

$$g_{j_k}(\boldsymbol{v}) = \frac{\boldsymbol{I}(Q_{j_k}, R_{j_k}) - \boldsymbol{I}(Q_{j_{k+1}}, R_{j_{k+1}})}{\sum_{m=1}^{N_g} \left|\boldsymbol{I}(Q_{j_m}, R_{j_m}) - \boldsymbol{I}(Q_{j_{m+1}}, R_{j_{m+1}})\right|}, \quad k = 1, \ldots, N_g. \quad (25)$$

The gradient of the normalized profile in (25) results in a lengthy but straightforwardly implemented result. The final expression still requires one to calculate the gradient of the image along the profile points—that is, $(\partial/\partial v_n)\boldsymbol{I}(Q_{j_k}, R_{j_k})$.

Applying the chain rule, we can write this partial derivative as follows:

$$\frac{\partial}{\partial v_n}\boldsymbol{I}(Q_{j_k}, R_{j_k}) = \frac{\partial}{\partial Q_{j_k}}\boldsymbol{I}\frac{\partial}{\partial v_n}Q_{j_k} + \frac{\partial}{\partial R_{j_k}}\boldsymbol{I}\frac{\partial}{\partial v_n}R_{j_k} \quad (26)$$

We do not have an analytical form of the image, $\boldsymbol{I}$, so we calculate its partial derivatives via a finite-difference operation. Also, we can write the partials of $\boldsymbol{Q}_j$ and $\boldsymbol{R}_j$ as follows:

$$\frac{\partial}{\partial v_n}Q_{j_k} = \frac{\partial}{\partial v_n}x_j + \left(\frac{(N_g-1)}{2} - k + 1\right)\sin\alpha_j \frac{\partial}{\partial v_n}\alpha_j \quad (27)$$

$$\frac{\partial}{\partial v_n}R_{j_k} = \frac{\partial}{\partial v_n}y_j - \left(\frac{(N_g-1)}{2} - k + 1\right)\cos\alpha_j \frac{\partial}{\partial v_n}\alpha_j. \quad (28)$$

The partials of x_j and y_j (the scaled and rotated coordinates of the jth LP) and α_j (the angle of the jth profile) have a straightforward analytic solution. This completes the calculation of the gradient terms in the objective function, J. Gradient descent can now be used to optimize the objective function to find the best boundary position and shape.

One important item to note about this formulation is the very small number of parameters that must be tuned for a given application. The length of the gray-level profiles, N_g, and the step-size used during the gradient descent optimization are the only two parameters that may need to be adjusted. This may be compared to the ASM implementation where up to five different parameters are typically adjusted for a new application. Section III presents some results using PSAM.

III. Experimental Method

PSAM requires a training set to incorporate shape and gray-level information about the object(s) of interest. To build the training set, a population of 30 mice were used. Of these 30 mice, 28 were healthy and two of them were afflicted with PKD. Only two PKD samples were available for this preclinical study because of the difficulties associated with housing and caring for these diseased animals. As described in the introduction, the mice were injected with contrast agent and three-dimensional (3-D) volumes of each animal were subsequently generated using an X-ray micro-CT scanner. Transaxial slices through the kidneys within the reconstructed volumes were used to train and test the PSAM algorithm. Several transaxial slices (from 1 to 6) containing kidneys were selected from each of the 30 subjects to create a database of 103 total kidney images. In the cases for which several slices were selected from the same animal, there is significant anatomic variation from slice to slice in both the kidneys and the structures surrounding the kidneys. As illustrated in the scout image in Fig. 2, all of the slices were manually selected using the L3 and L4 vertebrae as skeletal reference markers. In this region both the left and right kidneys are visible in an axial slice. Although beyond the scope of this paper, this information could be used to automate the process of selecting the kidney position.

To provide data for use in both the training set and as ground truth data for quantifying PSAM performance, the kidneys and spine were manually delineated in all 103 images by placing a collection of landmark points on the spine and kidney boundaries. Because of the flexibility of PSAM to handle both open and closed boundaries, the spine (an open boundary) was added to the two kidneys (closed boundaries) to serve as an anchor point because of its unique density in abdominal images. These manual segmentations were validated by a veterinarian radiologist. As is the case with similar PDM techniques, it is important for PSAM that the same number of LPs be used to delineate the boundary of each object (or objects) within the training set. Furthermore, the LPs must maintain the property of correspondence in that each LP must be placed in a similar position on the boundary of every object within the training set. This is typically accomplished by placing corresponding LPs on critical boundary features (e.g., corners), and then interpolating between these to create additional LPs that also lie on the boundary. In the case of the images being studied here, the kidneys in particular do not have consistent, distinguishing features (e.g., corners) to be used as "critical" features to help maintain correspondence. In addition, the varied kidney orientations make it challenging to consistently place LPs in the same position on the boundary within the training images. In this situation, it helps to manually add reference marks to the image that help guide the LP placement process. Fig. 3 shows a magnified kidney CT image that illustrates the process used to consistently label LPs along the kidney images. First, a single line was drawn that passes through the approximate center of both kidneys to divide the kidneys into upper and lower halves. Next, two additional lines were drawn that pass through the approximate center of each kidney and are perpendicular to the first line. As illustrated, the intersection points of these three lines with the kidney boundaries were used as the locations of the critical points. Two additional interpolating boundary points were then placed between these critical LPs to further delineate the kidney boundaries. Landmarking of the spine was carried out in a more consistent fashion in that three critical points were placed that mark the upper-left, upper-right, and bottom of the vertebra. Two additional LPs were placed on each side of the spine along the vertebra boundary.

IV. Results

To create the PSAM training set, one slice (and its corresponding LPs) was selected from the CT volume of 29 of the 30

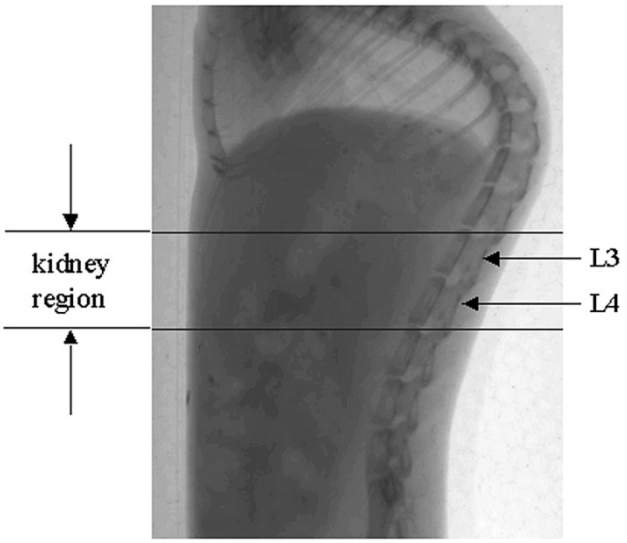

Fig. 2. Scout image showing the abdominal region from which the transaxial kidneys slices were extracted.

subjects. Prior to segmenting each of the 103 test images, PSAM was retrained (see Section II) using a set of 29 slices excluding the animal from which the current test image came. Regarding parameter settings, the gray-level profile length N_g was set to 7, the gradient step-size was chosen to be 3e-4, and these values were not changed during the entire testing process. The accuracy of segmentation was measured on each of the 103 images by comparing the PSAM segmentation to the manual segmentation of the corresponding image in the training set. The accuracy was quantified by measuring the average euclidean distance between the PSAM-generated LP positions and the position of the manually-placed LPs in the corresponding image. Three representative examples of the segmentation results are shown in Fig. 4. For the 103 cases, the average error in each LP position (i.e., the distance between the corresponding PSAM- and manually-generated LP) over all LPs was 2.4 pixels with a standard deviation of 1.4. The minimum and maximum average PSAM errors over all cases were 1.1 pixels and 5.8 pixels, respectively. The error measurement used here is quite conservative in that the error increases even if an LP moves along the correct boundary but away from its corresponding manually selected LP. Depending on the goal of the application, an error based on the distance from an LP to the true boundary (rather than the corresponding LP) may be more appropriate. This average error typically results in a kidney segmentation that is quite reliable in terms of its utility in screening for PKD. Fig. 5 shows a comparison of a PSAM result to the manually labeled boundary for the same image. The PSAM average error for this mouse was 2.3 pixels per LP. Note that there are minor differences in boundary position, but the kidneys are well delineated.

For comparison, conventional ASM was implemented as described in [2] and [6] and then applied to all 103 cases. After tuning multiple ASM parameters to achieve the best possible segmentation results, the average LP error was measured as 3.3 pixels with a standard deviation of 2.3. The minimum and maximum average ASM errors over all cases were 2.4 pixels and 8.0

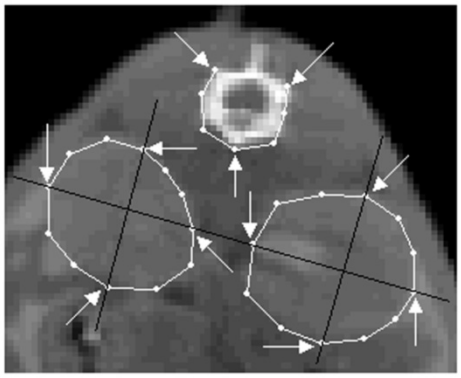

Fig. 3. Illustration of procedure to maintain LP correspondence on kidney boundaries across all images in the training set. Manually drawn lines through the kidney centers provide a reference for critical LP positions. The white arrows indicate the position of the critical LPs for the kidney and spine boundaries.

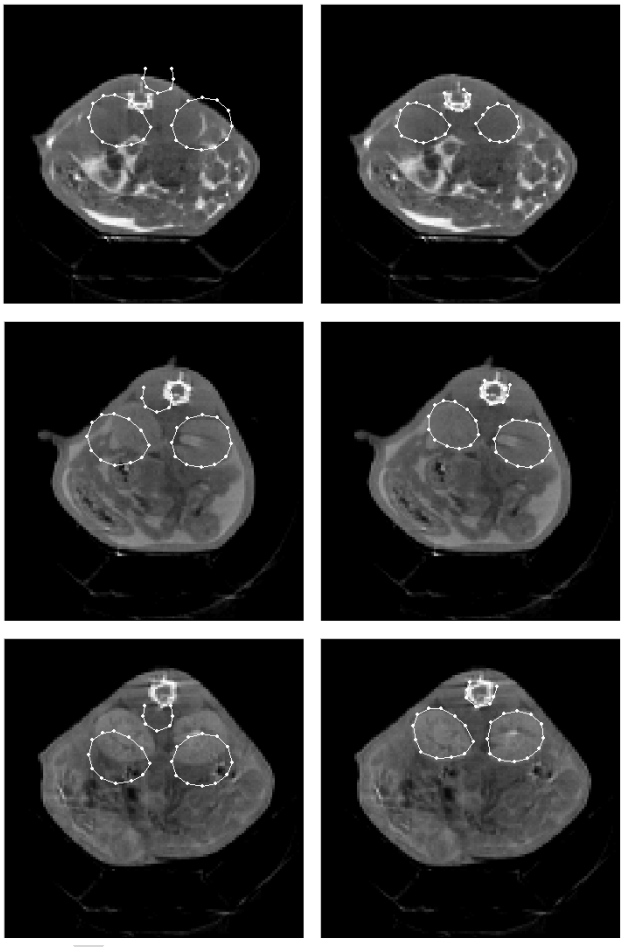

Fig. 4. PSAM results on three mouse kidney CT images. The images on the left are the initial PSAM position, and the images on the right show the final position of PSAM after convergence.

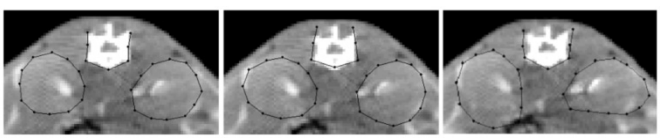

Fig. 5. Manually segmented image (left) compared to PSAM segmented image (center) and the ASM result (right). The calculated average LP position error for this case is 2.3 pixels for PSAM and 3.5 pixels for ASM.

pixels, respectively. This larger error led to significantly poorer kidney boundary delineation as can be seen in Fig. 5.

For all 103 images, an MOLV texture measurement was applied within the kidney boundaries to screen for the presence of PKD. The first step in the MOLV measurement is calculating the local variance of the gray-level values in the neighborhood (4×4 area) of each pixel within the kidney boundary. In the second and final step, the mean of the local variances across all of these pixel neighborhoods within each kidney boundary is computed. This mean can then be interpreted as a relative measurement of a region's texture. A mathematical expression for MOLV in one region (i.e., kidney) is

$$MOLV = \frac{1}{n(m-1)} \sum_{i=1}^{n} \left[\sum_{\forall j \in N(i)} \boldsymbol{I}_j^2 - \frac{1}{m} \left(\sum_{\forall j \in N(i)} \boldsymbol{I}_j \right)^2 \right] \tag{29}$$

where n is the total number of pixels contained within the kidney boundary, $N(i)$ is a set of pixels representing the local neighborhood around the ith pixel within the boundary, m is the number of pixels in $N(i)$ (note that $m = 16$ for a 4×4 local neighborhood), and I_j is the gray-scale value of the jth pixel in the image, I.

In all 103 cases, the presence (or absence) of PKD was correctly determined by setting the MOLV threshold level at 50, below which a kidney was flagged as healthy and above which a kidney was flagged as having PKD. Four representative texture measurements (two normal plus two PKD) from the 103 tested images are shown in Table I. The images from which these four texture measurements were extracted are shown in Fig. 6.

A repeatability experiment was performed by scanning two of the 30 mice a second time. In between scans, the animal was removed from the scanner and then placed back on the CT scanner bed. This process changed the position and orientation of the kidneys within the scanner and, hence, within the reconstructed CT volume. Six transaxial slices containing kidneys were segmented using PSAM and the accuracy was compared to the PSAM segmentation applied to the first scan of each animal. For the first subject the average LP position errors for the first and second scans were 2.8 and 2.5, respectively. For the second mouse, the same error measurements were calculated to be 1.8 and 2.8 pixels. These errors measured on the repositioned animal are well within the expected errors based on the tests performed on the original 103 images.

The segmentation errors that are encountered using the PSAM algorithm are generated from a couple of different sources. The two most likely sources are: 1) small inaccuracies in LP placement during PSAM training, and 2) PSAM boundaries getting trapped in local minima during the optimization process. The first source of error will diminish as the quality of the labeled training sets improves with operator experience. The second source of error may be reduced by trying optimization approaches other than gradient descent (e.g., simulated annealing).

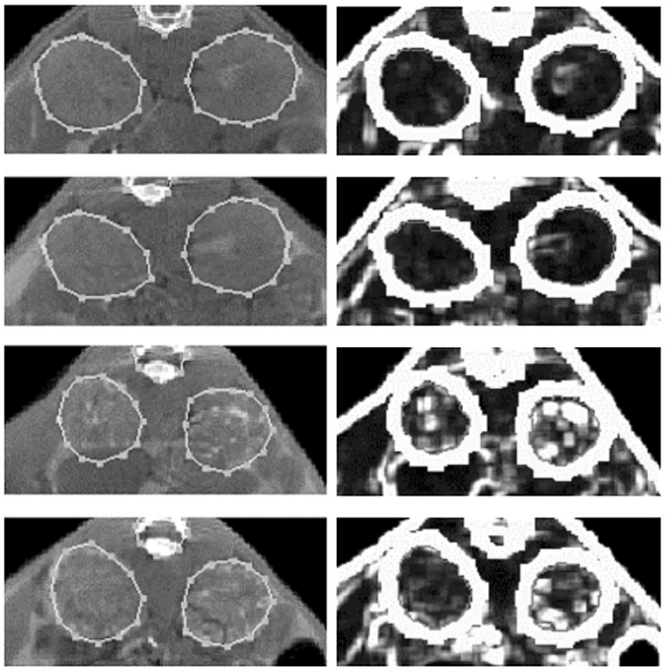

Fig. 6. Examples of two healthy kidney segmentations (rows 1 and 2) and two PKD segmentations (rows 3 and 4). The original images are on the left, and an 8-bit representation of the MOLV texture images are on the right.

TABLE I
REPRESENTATIVE MOLV KIDNEY TEXTURE MEASUREMENTS.

Normal Subject #	Texture Left	Texture Right
1	28	38
2	30	34
PKD Subject #	**Texture Left**	**Texture Right**
1	104	215
2	65	131

V. CONCLUSION AND FUTURE WORK

A theoretical formulation for a new statistical based deformable model algorithm (PSAM) has been presented that has several improvements over similar PDM-based techniques. Most notable is the formulation of an objective function that allows simultaneous optimization of the GSM and GLM. In the presence of missing and/or broken edge information, PSAM effectively uses the a priori shape and gray-level data gathered from the training set to create a model of gray-level appearance and a model of global shape that are combined into a single objective function. This unification of the models during optimization has several advantages in that it provides a general framework into which other (existing and future) PDM techniques fit well. Also, the number of parameters (only two) that need to be adjusted to achieve acceptable segmentation results is small relative to comparable techniques. Finally, as the results in Section IV demonstrate, the performance of PSAM is improved. This is particularly true for applications in

which the true object boundaries may be faint and/or missing. In these types of applications, simultaneous optimization of the gray-level and shape models becomes important to prevent some LPs from being shifted to stronger, but erroneous edges (see Fig. 5).

The PSAM method has been demonstrated as a very effective tool for segmenting the spine and kidneys in X-ray CT images of mice for the purposes of screening for PKD. Even in the few cases where a kidney boundary was partially missed, enough of the kidney was detected to allow accurate screening for PKD. Although not reported here, PSAM has been applied successfully to the segmentation of mouse heart and lungs as well as to mouse brain segmentation within X-ray CT images [11].

A three-dimensional version of the ASM algorithm has been developed for segmentation of a 3-D object embedded within a volumetric data set. One of the main challenges in the 3-D application of any PDM-based segmentation approach is the practical matter of manual training by an anatomy expert who may not be accustomed to viewing volumetric data on a computer screen. A new training approach has been developed to overcome this difficulty by presenting the 3-D volumetric data in a unique way that simplifies the overall process [12]. The improved PSAM optimization approach reported here is being updated for use with this 3-D algorithm.

REFERENCES

[1] M. Paulus, H. Sari-Sarraf, S. Gleason, M. Bobrek, J. Hicks, D. Johnson, J. Behel, and L. Thompson, "A new X-ray computed tomography system for laboratory mouse imaging," *IEEE Trans. Nucl. Sci.*, pt. 2, vol. 46, pp. 558–64, Jun. 1999.

[2] T. Cootes, C. Taylor, D. Cooper, and J. Graham, "Active shape models — Their training and application," *Comput. Vision Image Understand.*, vol. 61, no. 1, pp. 38–59, Jan. 1995.

[3] Y. Wang and L. Staib, "Boundary finding with correspondence using statistical shape models," *IEEE Comput. Vision Patt. Recog.*, pp. 338–345, Jun. 1998.

[4] C. Kervrann and F. Heitz, "A hierarchical Markov modeling approach for the segmentation and tracking of deformable shapes," *Graph. Models Image Processing*, vol. 60, no. 3, pp. 173–195, May 1998.

[5] N. Duta and M. Sonka, "Segmentation and interpretation of MR brain images: An improved active shape model," *IEEE Trans. Med. Imag.*, vol. 17, pp. 1049–1062, Dec. 1998.

[6] S. Gleason, H. Sari-Sarraf, M. Paulus, D. Johnson, and M. Abidi, "Automatic screening of polycystic kidney disease in X-ray CT images of laboratory mice," in *Proc. SPIE Conf. Medical Imaging*, vol. 3979, Feb. 2000, pp. 837–846.

[7] T. Cootes, G. Edwards, and C. Taylor, "Active appearance models," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 23, pp. 681–685, Jun. 2001.

[8] R. Duda and P. Hart, *Pattern Classification and Scene Analysis*. New York: Wiley, 1973.

[9] S. Gleason, H. Sari-Sarraf, M. Paulus, D. Johnson, and M. Abidi, "Statistical-based deformable models with simultaneous optimization of object gray-level and shape characteristics," *Proc. IEEE Southwest Symp. Image Analysis Interpretation*, pp. 93–95, Apr. 2000.

[10] C. Goodall, "Procrustes methods in the statistical analysis of shape," *J. Royal Stat. Soc. B*, vol. 53, no. 2, pp. 285–339, 1991.

[11] S. Gleason, "Development of a unified probabilistic framework for segmentation and recognition of semirigid objects in complex backgrounds via deformable shape models," Ph.D. dissertation, Univ. Tenn., Knoxville, TN, 2001.

[12] M. Dickens, S. Gleason, and H. Sari-Sarraf, "Volumetric segmentation via 3-D active shape models," *Proc. IEEE Southwest Symp. Image Analysis Interpretation*, pp. 248–252, Apr. 2002.