Automatic screening of polycystic kidney disease in x-ray CT images of laboratory mice

Shaun S. Gleason[*a], Hamed Sari-Sarraf[b], Michael J. Paulus[a], Dabney K. Johnson[a], Mongi A. Abidi[c]

[a]Oak Ridge National Laboratory[**], P.O. Box 2008, Oak Ridge, TN 37831-6011

[b]Texas Tech University, Lubbock, TX 79409

[c]University of Tennessee, Knoxville, TN 37996-2100

ABSTRACT

This paper describes the application of a statistical-based deformable model algorithm to the segmentation of kidneys in x-ray computed tomography (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 polycystic kidney disease (PKD). The algorithm is based on active shape models (ASMs) initially developed by Cootes, *et. al.* Once the segmentation is complete, texture measurements are applied within kidney boundaries to detect the presence of PKD. The challenges associated with the segmentation of mouse kidneys (non-rigid organs) are presented, and the motivation for using ASMs in this application is discussed. Also, improvements were made to published ASM methods that may be generally helpful in other segmentation applications. In 15 of the 18 cases tested, the mouse kidneys and spine were detected with only minor errors in boundary position. In the remaining three cases, small parts of the kidneys were missed and/or some extra abdominal tissue was inadvertently included by the boundary. In all 18 cases, however, the kidneys were successfully detected at a level where PKD could be automatically screened for using mean-of-local-variance (MOLV) texture measurements.

Keywords: Computer-aided diagnosis, deformable models, active shape models, medical image segmentation, computer vision

1. INTRODUCTION

The Oak Ridge National Laboratory (ORNL) has a world class 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). A Laboratory-Directed Research and Development (LDRD) program was recently completed at ORNL with the objective of accelerating the process of screening these mice for internal abnormalities. A new high-resolution, x-ray CT instrument called MicroCAT has been developed as a part of this new program[1].

The 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 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. The algorithm is based on ASM developed by Cootes, *et. al*[2,3]. Once the segmentation is complete, MOLV texture measurements are applied within kidney boundaries to detect the presence of PKD.

[*]Correspondence: Email: gleasonss@ornl.gov; WWW: http://www-ismv.ic.ornl.gov/~gleason; Phone: 423 574 8259; Fax: 423 574 6663.
[**]Managed by Lockheed Martin Energy Research Corporation for the U.S. Department of Energy under contract DE-AC05-96OR22464.

This paper first describes the challenges associated with the segmentation of mouse kidneys (and non-rigid organs, in general). Next, the concept behind deformable model-based segmentation algorithms is presented. It then goes on to present the motivation for using ASMs in this application. Also, the paper outlines the general theory behind ASMs and describes improvements that were made to published ASM methods that may be generally helpful in other segmentation applications. Finally, the results of a PKD screening experiment are presented.

2. MEDICAL IMAGING AND DEFORMABLE MODELS

The field of medical imaging is one where anatomical structures may be somewhat predictable in terms of appearance (e.g. organ shape and location), but variations are always encountered from subject-to-subject and within the same subject over time. Also, there are areas of the anatomy (e.g. the abdominal cavity) where the appearance, location, and background of structures are quite unpredictable, making the segmentation task even more intractable. The difficulties created by these variations in the organ and its background are, at least in part, why semi-rigid organ segmentation is not as well-studied as rigid organ (or rigidly-enclosed organ, i.e., brain) segmentation. Also, in many cases of medical image analysis, there is an abundance of *a priori* information available in the form of the patient's own historical records as well as imagery from a potentially large population of other patients that could be used for algorithm training. In summary, a useful medical image segmentation algorithm should be capable of handling problems with the following characteristics:

- semi-rigid objects,
- complex backgrounds,
- faint, obscured, or partially missing object boundaries, and
- available *a priori* information on object appearance.

As an example of a medical imaging application with these characteristics, consider the problem of detecting kidneys within the abdomen of laboratory mice. A set of kidney cross-sectional images from 12 different mice is shown in Fig. 1. These images were acquired using the aforementioned MicroCAT x-ray CT system. Note the variability in shape, intensity, position, 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 very faint boundaries around many of the kidneys (e.g. second row, middle image). One particular class of segmentation algorithms that is adept at handling these types of variabilities is deformable models, and they are briefly described next.

2.1. Deformable Models

Deformable models belong to a class of algorithms that typically use a boundary representation (2D deformable contour) or a surface representation (3D deformable surface) of an object for segmentation purposes. The deformable contour is typically initialized by placing a contour on the image, and then its position and shape are iteratively adjusted to best fit the object's boundary. The characteristics of the contour (e.g. local roughness, overall shape) depend on a set of parameters that is adjusted and tuned via an optimization process to best fit the information about the object as represented by the image data. These parameters that control the general form of the resultant contour are typically referred to as *internal* constraints. The image data generate forces, referred to as *external* forces, that attract or repel the model. In their foundational work on active contours, or snakes, Kass *et. al*[4] formulated the internal constraints and external forces as internal and external energy terms in an objective function that, when minimized, corresponds to the best fit of the model to the image data.

2.2. Active Shape Models

Within the class of deformable models, ASM is one of the algorithms best-suited for applications where *a priori* information is available for incorporation into the segmentation process. ASM was first introduced by Cootes, *et. al*[2], and more recent research on ASM-based statistical deformable models has been presented by Wang, *et. al*[5], Kervrann *et. al*[6], and Duta *et. al*[7]. ASM uses a training set of images that have been manually segmented via the placement of a collection of landmark points (LPs) along the boundary of an object. Information is extracted from this training data to create two distinct models, a shape-model (SM) and a gray-level model (GLM), that together constitute the overall model in this hybrid segmentation approach. SM is created using the spatial location and the interrelationship of LPs across the training images, while GLM is created using

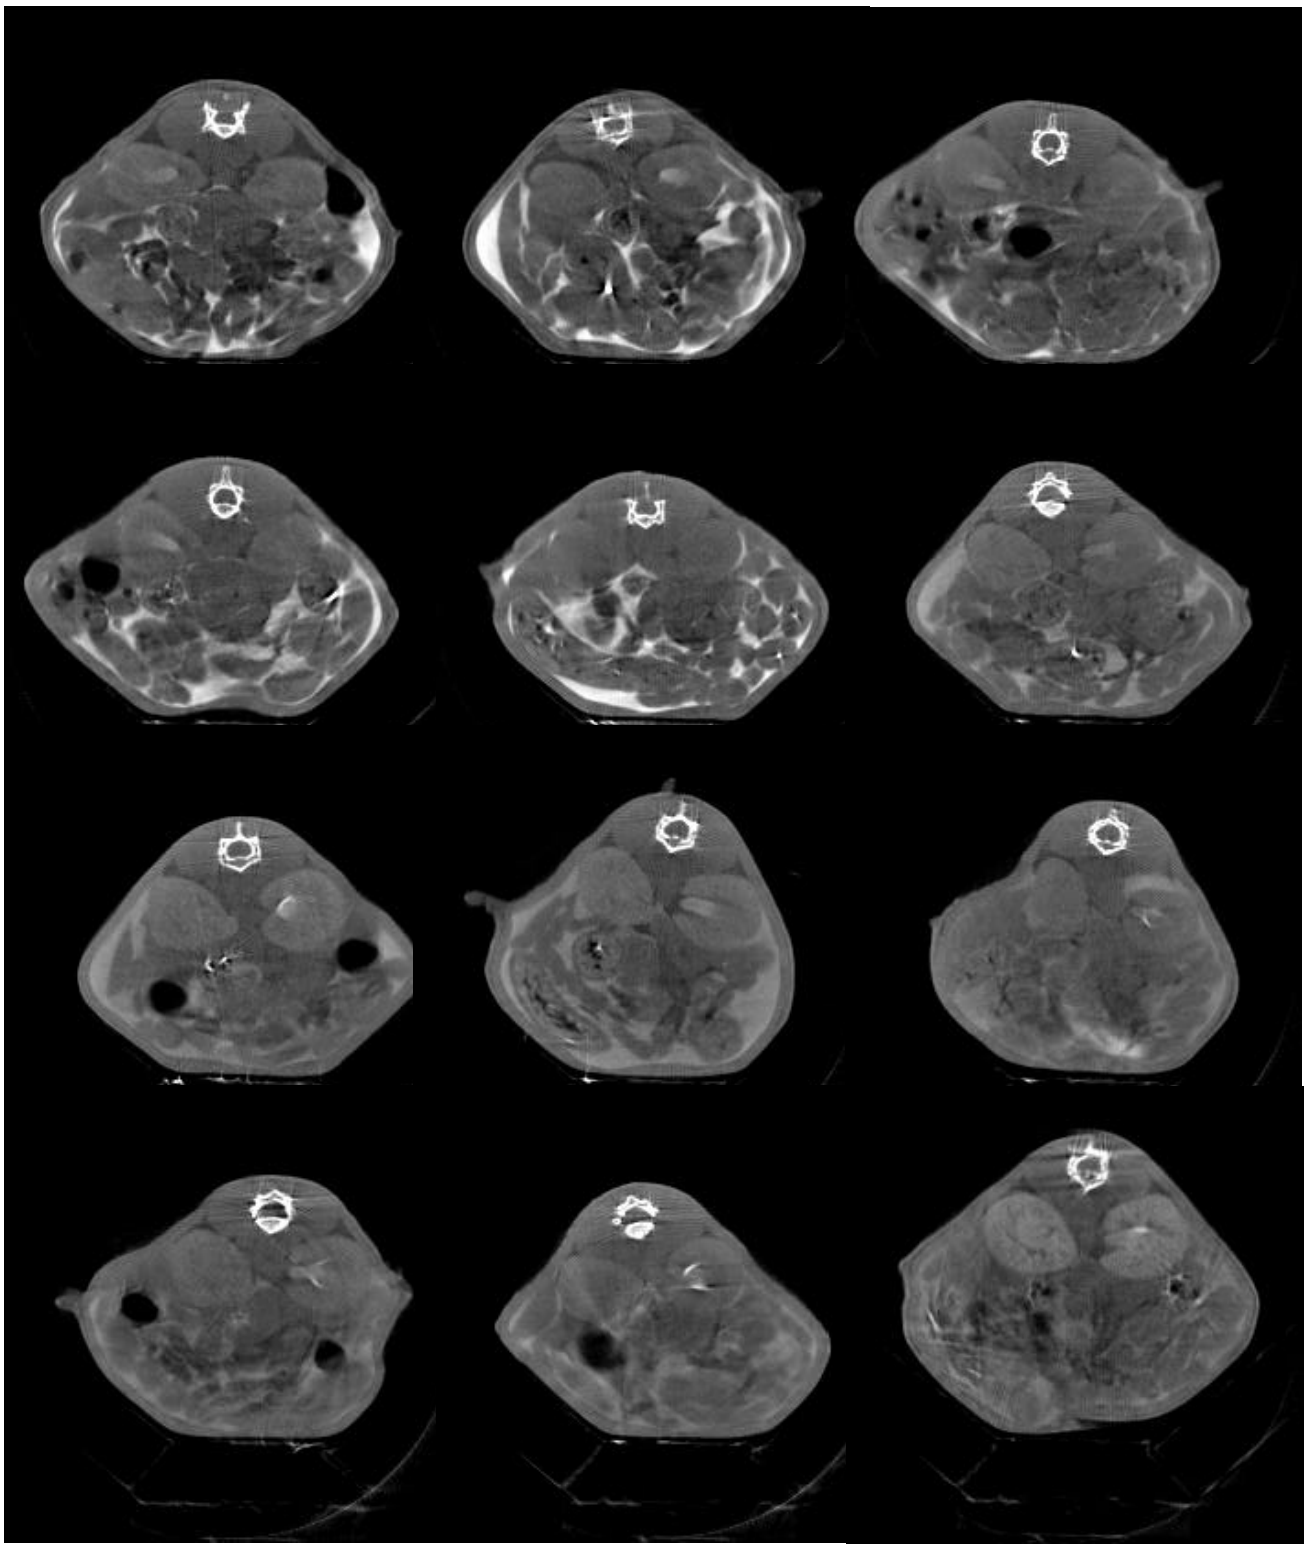

Figure 1. X-ray CT abdominal cross-sections of 12 different mice. The white structure at the top of each image is the spinal column. The two elliptical shaped structures to the lower-left and right of the spinal column are the kidneys. The black spots in the abdomen are air pockets and the distributed white abdominal structure (especially the first 5 images) is excess contrast agent.

local gray-scale gradient characteristics of sampled profiles within the training images. ASM is a multiresolution search technique, so GLMs and SMs are created during training (and applied during testing) at each of the predefined resolutions.

To put the optimization of ASM in the general context of deformable contours described earlier, SM imposes the *internal* shape constraints on the final solution, and GLM is influenced by the *external* forces created by the image data. ASM is initialized by placing an "average" contour on a test image. The shape and position of this initial contour are then iteratively updated to optimize their fit to the underlying image information. As in other deformable contour algorithms, the optimization of ASM achieves a compromise between the internal constraints imposed by SM, and the external forces influencing GLM. An iterative scheme is employed whereby GLM finds a new set of LPs based on the best fit to the image data and passes that LP set to SM. SM then constrains the shape of those LPs to fall within the range of allowable shapes as defined by the training set, and passes this constrained set of LPs back to GLM. This process repeats until the constrained set of LPs generated by SM stops changing. ASM is an excellent starting point to solve the problem of screening for PKD in laboratory mice. ASM does make substantial use of *a priori* information through the creation of the GLM and SM components. The knowledge built into these components during the training process allows ASM to segment objects with faint and/or missing boundaries with remarkable accuracy. Another important advantage that the ASM approach has over other deformable model algorithms is its intrinsic ability to simultaneously handle multiple boundaries, both open and closed.

3. IMPROVED STATISTICAL SHAPE MODEL

Although an excellent starting point for the problem at hand, ASM does suffer from shortcomings that had to be solved before it could be effective in this fully automated screening application. During the ASM implementation phase, two improvements were made to the published algorithm that positively affected the performance of both the SM and GLM components of ASM. First, noting that ASM is a multi-resolution technique, resolution-dependence was built into the SM to provide increased shape constraints at lower image resolutions, and more shape flexibility at higher resolutions. Second, GLM was modified to allow incorporation of both derivative and absolute-intensity information while updating the ASM boundary positions. This section describes the two extensions that were made to the original ASM formulation.

3.1. Resolution-Dependent Shape Model

The first change in the ASM approach was made to the SM constraint used to control the final boundary shape. The SM constraint given by[3]

$$-s\sqrt{\lambda_i} \le b_i \le s\sqrt{\lambda_i},\ \forall i\ =\ 1, \ldots, t \tag{1}$$

is applied to the final shape parameter vector, $\boldsymbol{b}$, to constrain shapes that vary significantly from those in the training set. The variable λ_i is the ith eigenvalue of the shape parameter covariance matrix generated by the training samples. The challenge is to select a value of s (the number of standard deviations away from the mean that each shape parameter is allowed to vary) that is small enough to keep the final shape within the bounds defined by the training set, yet large enough to allow the boundary to deform to fit the semi-rigid object in the image.

Typically, the SM shape constraint parameter, s, is held constant for each resolution of the multiresolution search. It was discovered that this can lead to undesirable segmentation results, especially in the low-resolution images. Because the pixels are relatively large in the low-resolution images, an error in the GLM fit of only a pixel or two can result in a large variation in shape. As an example, see the images in the left-hand column of Fig. 2. The GLM has a difficult time with the low-resolution image because of some competing edges below the right kidney. Note how the shape is deformed (right kidney larger and lower than the left), and this error propagates through the mid- and full-resolution images. The final result (bottom-left, Fig. 2) has missed the right kidney as well as the spine.

To prevent this significant shape deformation at low-resolution, a resolution-dependent scaling term, r, was added to the original SM equation, Eqn. (1), as follows

$$-rs\sqrt{\lambda_i} \le b_i \le rs\sqrt{\lambda_i},\ \forall i\ =\ 1, \ldots, t, \tag{2}$$

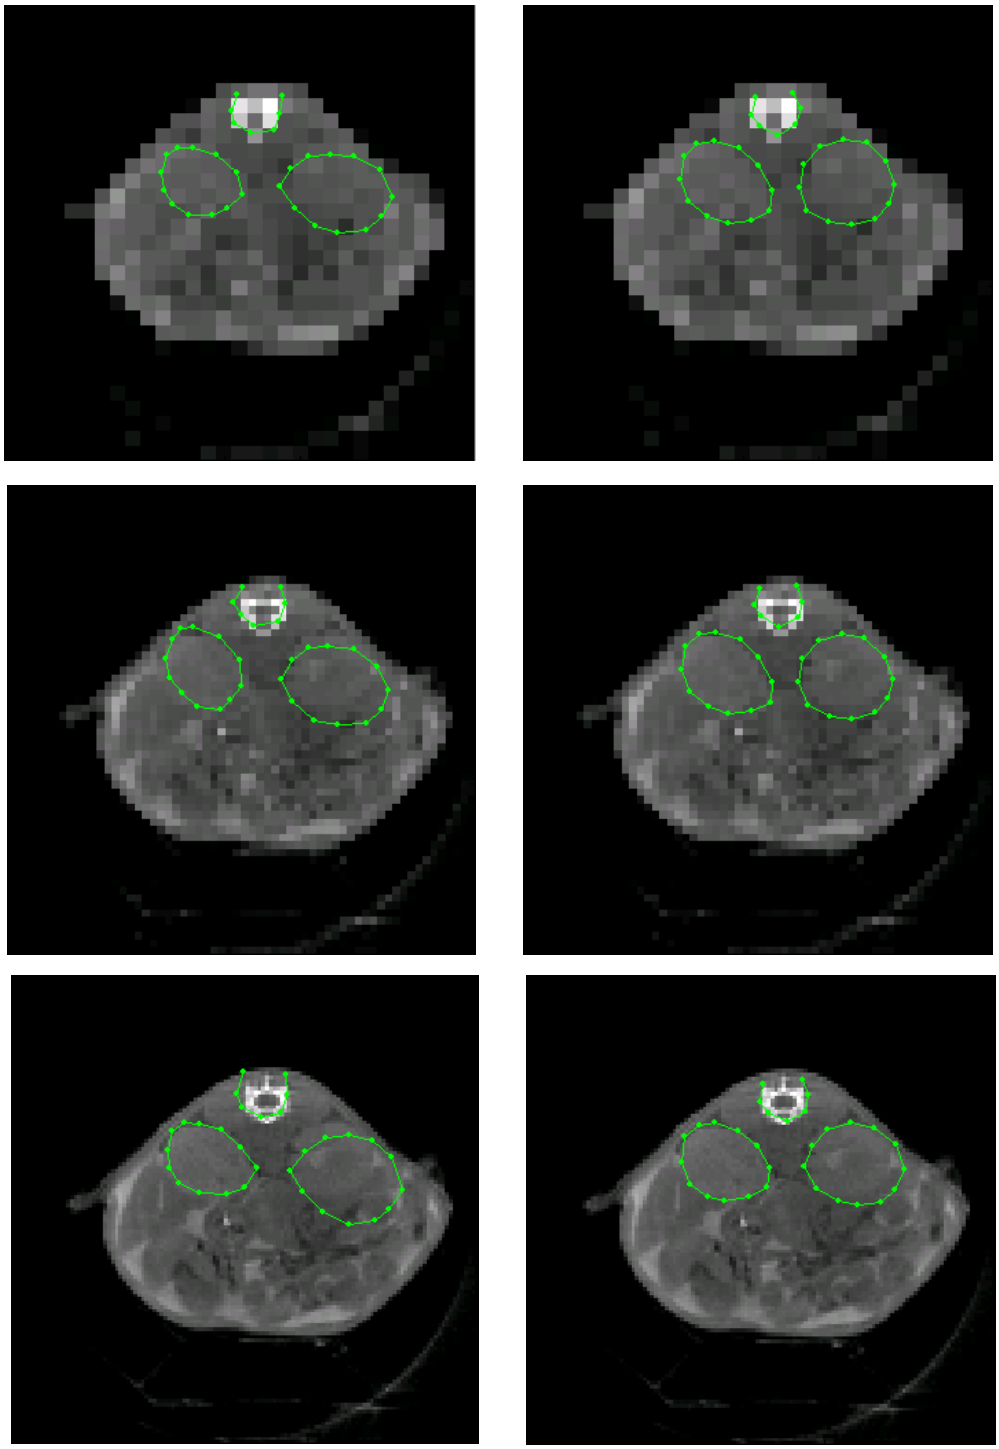

Figure 2. Example of multiresolution ASM search results with standard, resolution-independent shape constraint (left column) and with resolution-dependent shape constraint (right column). Images are 1/4-resolution (top), 1/2-resolution (middle), and full-resolution (bottom).

where r is equal to the image resolution (e.g. 0.25, 0.5, or 1.0 for Fig. 2). This tightens the constraint on the shape at low resolution, and leads to the results shown in the right-hand column of Fig. 2. This same approach has also been applied with successful results to the ellipsoid-constraint[3] as follows

$$\sum_{i=1}^{t} \frac{b_i^2}{\lambda_i} \leq rM_t, \tag{3}$$

where M_t is the Mahalanobis distance threshold from the current shape estimate to the mean shape and t is the number of parameters in the shape vector, $\boldsymbol{b}$. Because of the inherent instability of GLM at low resolution, a more strict shape constraint (box-constraint, Eqn. (2), or ellipsoid-constraint, Eqn. (3)) leads to better overall segmentation results.

3.2. Extended GLM

One of the difficulties in many medical image segmentation tasks is the possibility of a complex background. This complication is especially problematic when detecting structures in the abdomen of the subject because of the variability in abdominal structure (see Fig. 1). Such a background can in many cases generate spurious edge information that will confound most segmentation techniques, including ASM. Because the GLM component relies on gradient profile information, spurious edges can generate multiple possible matches, and, hence the possibility of an erroneous boundary position. If the shape of this erroneous boundary happens to fall within the acceptable global shape constraints as indicated by the training set, SM has no chance of correcting the erroneous decision. This shortfall can be attributed to at least one problem with the ASM approach: a reliance on only image gradient information during GLM optimization. Absolute intensity information from the image has also shown to valuable (along with gradient information) when attempting to locate the object boundaries.

To provide a means of encompassing both gradient and absolute intensity information into GLM, an additional profile is captured during GLM training. This second profile is also composed of sampled intensity values along the normal line to the LP, but differs from the original profile in two important ways. First the absolute intensities (not the derivatives) are used when building the GLM, and second, the image is only sampled on the *inside* of the object boundary. The reason for this is that in many segmentation applications, including the ones presented here, the object has a more predictable appearance than its background. Hence, the gray-level appearance of the object (excluding the background) can be effectively modeled by sampling only the interior of the object's boundary. It should be noted that Cootes *et. al* have published work on active appearance models (AAMs)[8] that incorporate absolute gray-level intensities of objects in a quite different manner. The AAM actually uses the values of *all* pixel intensity values within an object during training, not just profiles. Also, the AAM only uses gray-level information, not a combination of gray-level and derivative information as presented here.

The GLM objective function, $f(\boldsymbol{g}_d)$, was modified to include this new information in the new GLM objective function

$$f(\boldsymbol{g}_d, \boldsymbol{g}_i) = \alpha f_d(\boldsymbol{g}_d) + \beta f_i(\boldsymbol{g}_i), \tag{4}$$

where α and β are constants (range 0 to 1) that determine the relative contribution of the derivative- and intensity-model terms, $\boldsymbol{g}_d$ is the gradient profile, $\boldsymbol{g}_i$ is the new absolute intensity profile, and f_d, f_i are Mahalanobis distances defined as:

$$f_d(\boldsymbol{g}_d) = (\boldsymbol{g}_d - \overline{\boldsymbol{g}_d})^T \boldsymbol{S}_d^{-1} (\boldsymbol{g}_d - \overline{\boldsymbol{g}_d}), \tag{5}$$

$$f_i(\boldsymbol{g}_i) = (\boldsymbol{g}_i - \overline{\boldsymbol{g}_i})^T \boldsymbol{S}_i^{-1} (\boldsymbol{g}_i - \overline{\boldsymbol{g}_i}). \tag{6}$$

In these equations, $\overline{\boldsymbol{g}_d}$ and $\overline{\boldsymbol{g}_i}$ are the average gradient and absolute intensity profiles for a given LP across the training set, and $\boldsymbol{S}_d$, $\boldsymbol{S}_i$ are the covariance matrices for the gradient and absolute intensity profiles, respectively. An example of an improved result using this updated GLM is shown in Fig. 3. The spurious edge generated by the small pocket of contrast agent near the right kidney pulled the final boundary away from the true kidney boundary (segmentation result on left). This problem was corrected using the enhanced GLM (segmentation result on right).

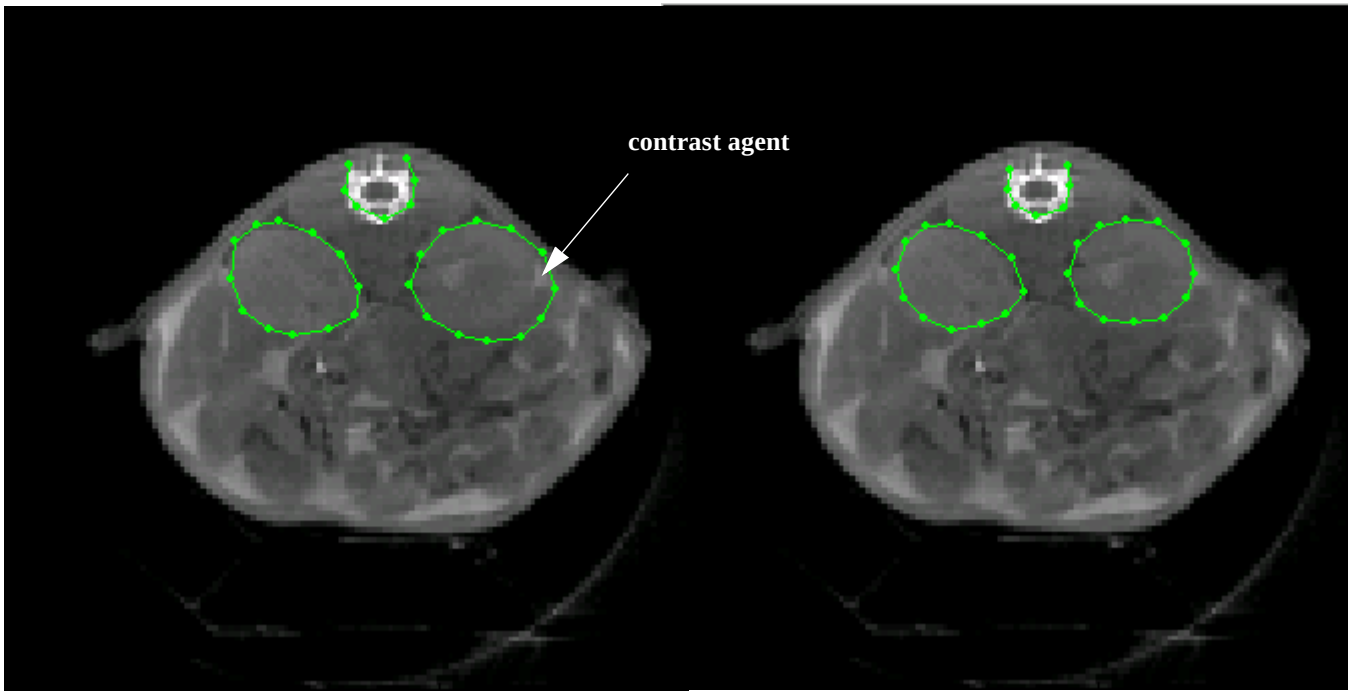

Figure 3. An example of improved segmentation using original GLM formulation (left) and extended GLM (right). Note the exclusion of the small spot of contrast agent next to the right kidney.

4. SCREENING RESULTS

Because ASM is a statistical-based, deformable shape-model, it requires a training set to incorporate shape and gray-level information about the object(s) of interest. To build a training set, a collection of 16 axial CT images of healthy mouse kidneys were captured using the micro-CT scanner at ORNL. The kidneys and spine were then manually segmented in each of the 16 images by placing a collection of landmark points on the spine- and kidney-borders. The enhanced ASM described in the previous section was trained and applied to each of the 16 images for the purpose of segmenting both kidneys and the spine of the mouse. Because of the flexibility of ASM to handle multiple boundaries, both open and closed, the spine (an open boundary) was added to the two kidneys (closed boundaries) to serve as an anchor point. Using a hold-one-out approach, the accuracy of segmentation was measured by comparing the ASM segmentation result to the corresponding manually-segmented image in the training set. Also, the shape-model was applied to two additional PKD images. For all 18 kidney images, an MOLV texture measurement was applied within the kidney boundaries to screen for the presence of PKD. Some of the segmentation results are shown in Fig. 4.

In 15 of the 18 cases, the mouse kidneys and spine were completely detected with only a few minor errors in boundary position. In the remaining three cases, small parts of the kidneys were missed and/or some extra abdominal tissue was inadvertently included by the shape-model. In all 18 cases, the kidneys were successfully detected at a level where PKD could be automatically screened for using the aforementioned texture measurement. In all 18 cases the presence (or absence) of PKD

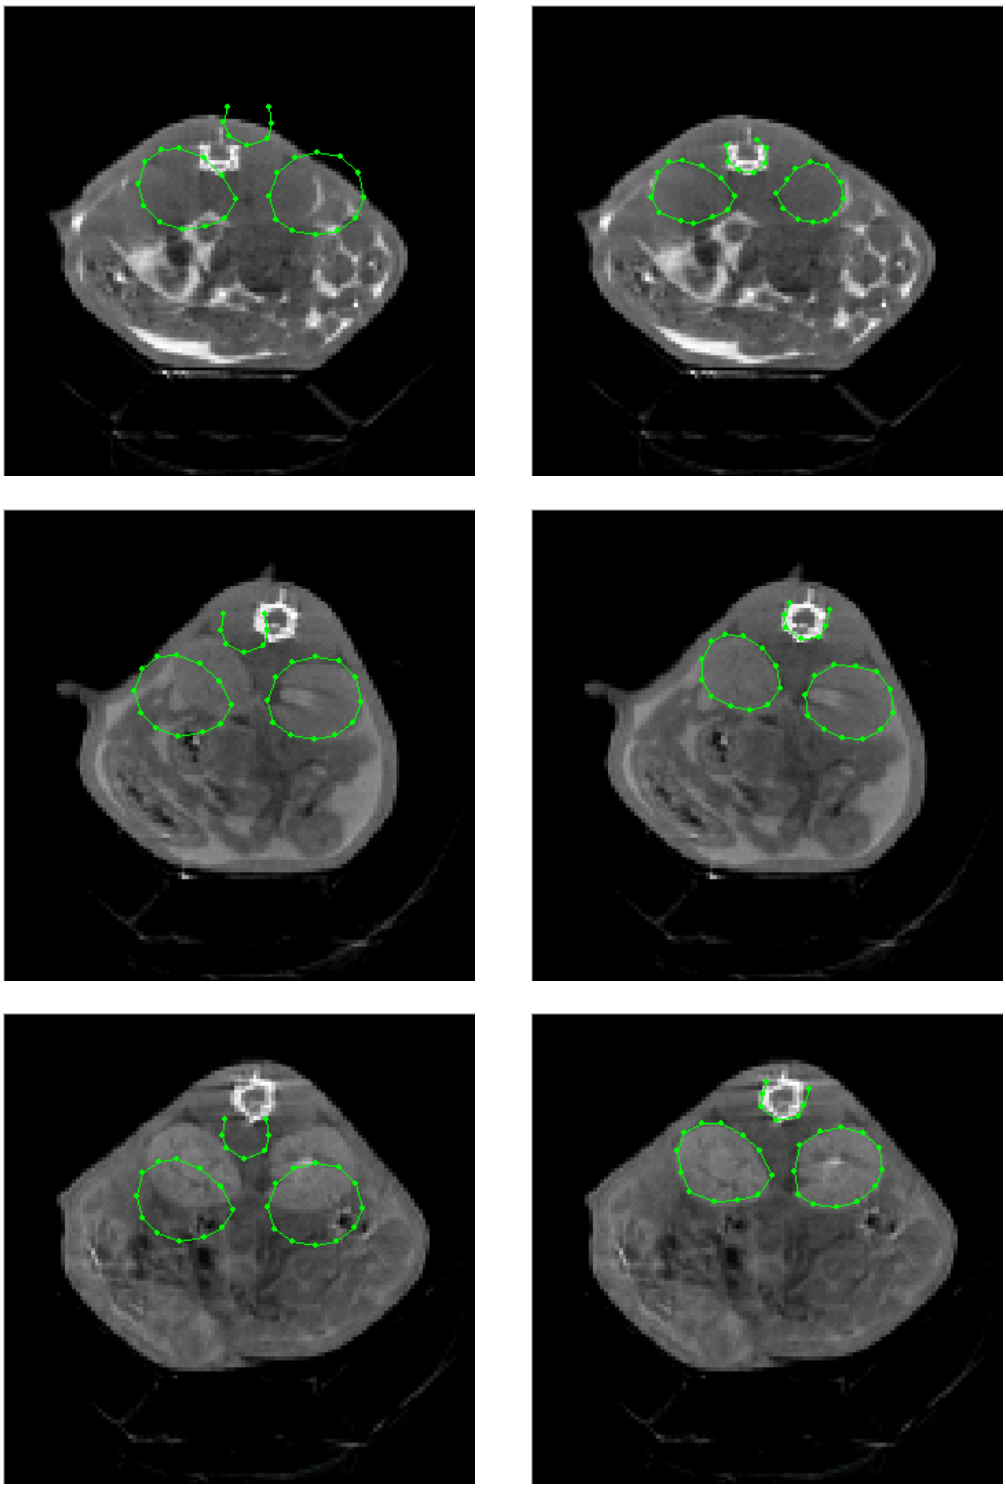

Figure 4. ASM results on three mouse kidney CT images. The images on the left are the initial ASM position, and the images on the right show the final position of the ASM after convergence.

was correctly determined. The texture measurements for 4 of the 18 tested images (two normal plus two PKD) are shown in Table 1. The images from which these four texture measurements were extracted are shown in Fig. 5.

Table 1: 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

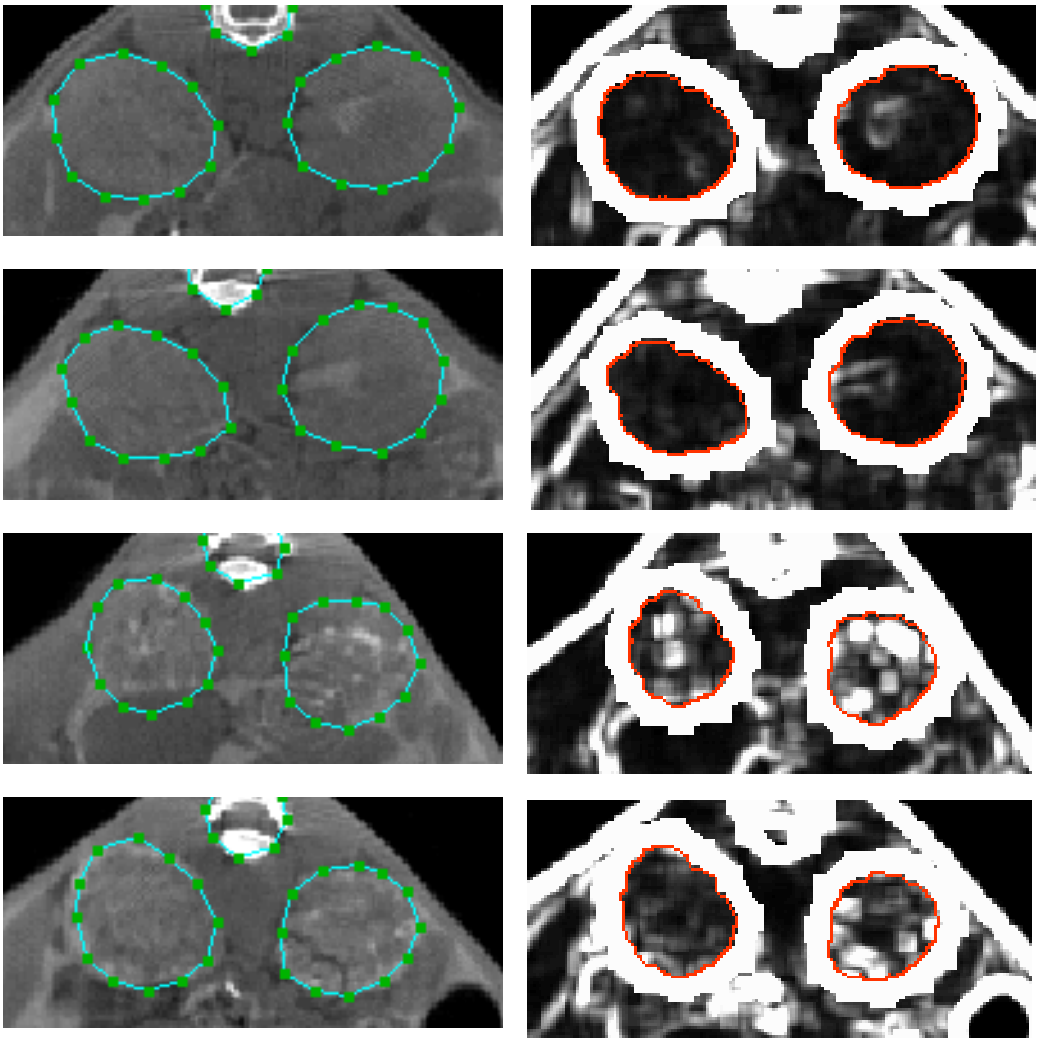

Figure 5. Examples of two healthy kidney segmentations (rows 1 & 2) and two PKD segmentations (rows 3 & 4). The original images are on the left, and the MOLV texture images are on the right.

5. CONCLUSIONS AND FUTURE WORK

The enhanced ASM method presented in this paper has been a very effective tool for segmenting the spine and kidneys in x-ray CT images of mice. In the presence of missing and/or broken edge information, this approach effectively uses the *a priori* shape and gray-level data gathered from the training set to segment the spine and kidneys. Even in the cases where a kidney was partially missed, enough of the kidney was detected to allow automatic screening for PKD. Although good performance was obtained using the enhanced ASM approach presented here, there are opportunities for further improvement.

Recall that the original ASM approach separately optimizes the SM and the GLM components of the ASM. To summarize, GLM searches the image gradient profiles and updates the LP positions with no regard as to how well they will fit the global constraints imposed by SM. Similarly, SM fits a new (model-generated) set of LPs to the GLM-suggested LPs with no regard as to how this new set matches the gradient profile information contained in the GLM. A reformulation of the ASM optimization approach is needed that uses the gray-level and global shape information captured in the training set in a more robust and efficient way that simultaneously considers the effects of both during the deformation process.

The lack of interaction between the GLM and SM creates another unfortunate situation in that no confidence measure is available that approximates how well the ASM segmentation algorithm has performed on any given image. In some fully automated applications, this renders the ASM ineffective. The previously mentioned reformulation of the optimization scheme should also provide a means to generate a performance confidence metric that will allow the use of this new statistical-based deformable shape model in a fully automated application.

REFERENCES

1. M. J. Paulus, H. Sari-Sarraf, S. S. Gleason, M. Bobrek, J. S. Hicks, D. K. Johnson, J. K. Behel, L. Thompson, "A New X-ray Computed Tomography System for Laboratory Mouse Imaging," *IEEE Trans. on Nuclear Science* **46**, no.3, pt.2, p. 558-64, June, 1999.
2. T. Cootes, C. Taylor, D. Cooper, and J. Graham, "Active shape models — their training and application", *Computer Vision and Image Understanding* **61**, no. 1, pp. 38-59, Jan. 1995.
3. T. Cootes and C. Taylor, "Active shape models", University of Manchester Technical Report, http:// www.wiau.man.ac.uk, 1998.
4. M. Kass, A. Witkin, D. Terzopoulos, "Snakes: Active Contour Models," *International Journal of Computer Vision*, 321-331, 1988.
5. Y. Wang, and L. Staib, "Boundary finding with correspondence using statistical shape models", *IEEE Proc. Computer Vision and Pattern Recognition*, pp. 338-345, 1998.
6. C. Kervrann, and F. Heitz, "A hierarchical markov modeling approach for the segmentation and tracking of deformable shapes", *Graphical Models and Image Processing* **60**, no. 3, pp. 173-195, May 1998.
7. N. Duta and M. Sonka, "Segmentation and interpretation of MR brain images using an improved knowledge based active shape model", *Proc. IPMI '97 Conf. Lecture Notes in Computer Science*, pp.375-380, 1997.
8. T. Cootes, G. Edwards, and C. Taylor, "Active appearance models", *Proc. European Conference on Computer Vision* **2**, pp 484-498, 1998.