

# Shape-Based Registration of Kidneys Across Differently Contrasted CT Scans

Fernando Flores-Mangas, Allan D. Jepson  
*Department of Computer Science*  
*University of Toronto*  
*Toronto, ON, Canada*  
{mangas, jepson}@cs.toronto.edu

Masoom A. Haider  
*Department of Medical Imaging*  
*Princess Margaret and Mount Sinai Hospitals*  
*Toronto, ON, Canada*  
m.haider@utoronto.ca

**Abstract**—We present a method to register kidneys from Computed Tomography (CT) scans with and without contrast enhancement. The method builds a patient-specific kidney shape model from the contrast enhanced image, and then matches it against automatically segmented candidate surfaces extracted from the pre-contrast image to find the alignment. Only the object of interest is used to drive the alignment, providing results that are robust to near-rigid relative motions of the kidney with respect to the surrounding tissues. Shape-based features are used, as opposed to intensity-based ones, and consequently the resulting registration is invariant to the inherent contrast variations. The contributions of this work are: a surface grouping and segmentation algorithm driven by smooth curvature constraints; and a framework to register image volumes under contrast variation, relative motion and local deformation with minimal user intervention. Encouraging experimental results with real patient images, all with various kinds and sizes of kidney lesions, validate the approach.

**Keywords**-3D Registration, surface segmentation, shape modeling.

## I. INTRODUCTION

Renal lesions are very common, and while most of them are benign cysts, some are renal cell carcinomas. Identification of the type of lesion requires measuring the brightness differences between two (routine) CT scans: one taken prior to the injection of a contrast agent (pre-contrast), and another one taken after injection (post-contrast). Currently, evaluating the level of contrast enhancement across corresponding image locations is a labor-intensive procedure, which requires a radiologist to visually inspect and manually match corresponding 3D points. Matching of regions is necessary because kidneys move between the pre- and post-contrast images due to respiration and patient's bulk motion over time. In addition, matching precision must be within the voxel size for the test to be informative for diagnosis.

This paper presents a shape-based registration method that enables determination of the level of contrast enhancement for the entire kidney, and does so with minimum user intervention. The approach builds a patient-specific 3D shape model of the kidney using the post-contrast image. This model is then matched against candidate surfaces extracted from the pre-contrast image to find the alignment. Candidate surfaces are automatically generated using a novel grouping

and segmentation approach, driven by iteratively redefined curvature constraints. The algorithm reaches a solution when several surfaces (built from the pre-contrast volume) agree on the location of the model (built from the post-contrast volume). Figure 1 shows a schematic overview of the algorithm. The approach overcomes the limitation of existing algorithms where all image structures influence the resulting registration transformation. Experimental Results with 20 datasets of actual patients with renal lesions indicate that our method efficiently (within seconds) produces highly accurate (within voxel size) rigid registration results, useful for rapid diagnosis.

## II. RELATED WORK

The large variety of registration methods available in the literature responds to the proportional number of imaging modalities, the nature and domain of the registration transformations and the specific needs and constraints of each application [1]. Twining et al. [2] present a method for group-wise registration, statistically modeling images using an information theory approach, where the coding length of a set of images is minimized, maximizing correspondence. While their method would non-rigidly co-register full abdominal volumes down to a canonical reference model, the proposed objective is not robust to contrast variations.

Statistical modeling is also exploited by Pizer et al. [3], but this time using medial representations to characterize the geometry of anatomical structures from a class of objects. Registration using this method requires segmentations for both the pre- and post-contrast image structures. Instead, our method requires segmentation from the post-contrast image only, where the kidney is most salient, and avoids segmentation on the pre-contrast image, where segmentation may be impractical due to lack of contrast. Segmentations are also needed by methods based on the Iterative Closest Point (ICP) algorithm [4], [21]. Alternatively, when provided with crude edge information, ICP methods suffer from a reduced basin of convergence to a sensible solution, so they must be carefully initialized. This is precisely one of the challenges addressed by the method presented here.

Mutual Information methods [5], [6] do take into account intensity variations across image sets, but they generally treat

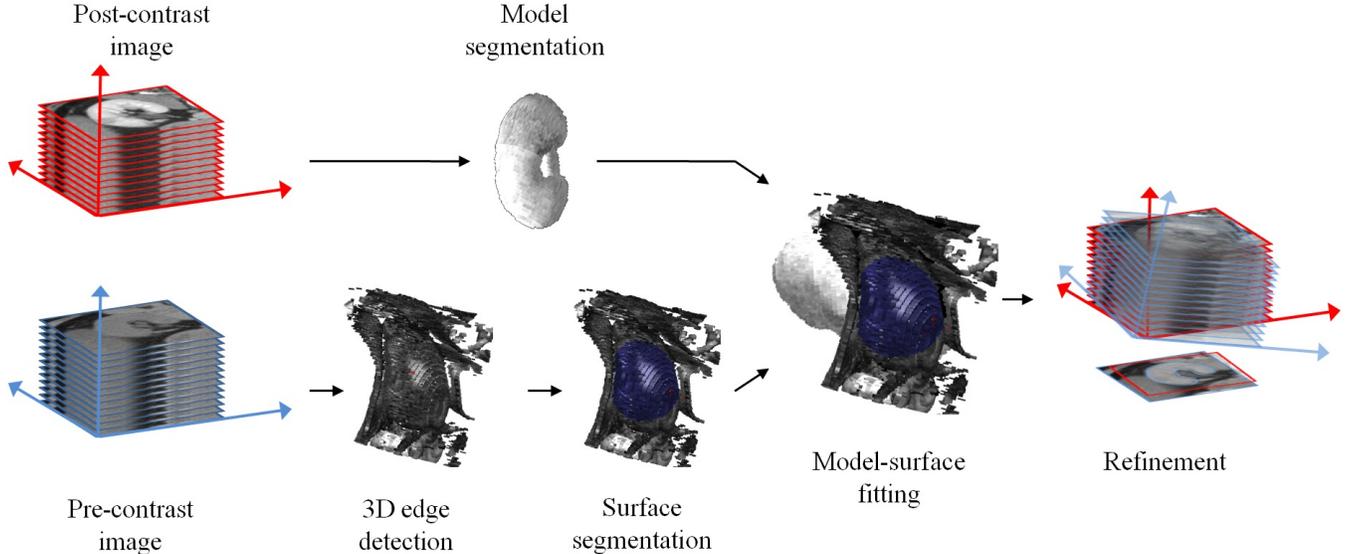


Figure 1: Algorithm overview. A 3D model is obtained from the post-contrast image. Surfaces are segmented from the pre-contrast image and aligned to the model. A refinement stage is applied last.

all image contents equally, making the result sensitive to relative motion. Sun et al. [7] acknowledge the effects of contrast variations and propose a contrast-invariant similarity metric but their method only works in 2D. Spherical Harmonic-based shape models have also been proved effective; in [8], the authors propose a sampling-based search algorithm that greedily minimizes surface distances using coordinate gradient descent on the rotation parameters and then jointly optimizing, but again, segmentations for both datasets are needed.

More recently, Heldmann et al. [9] proposed a method to non-rigidly align multi-phase CT images of liver. They point out similar difficulties (*i.e.*, relative motion, contrast variation), and their objective function includes an image gradient alignment term, which is in principle similar to our method, but instead we use image gradient information to segment coherent surfaces as more robust features for alignment. Conversely, Kwon et al. [10] extract local image features from both images using a 3D corner descriptor; corresponding features are then used to find the alignment, balancing transformation smoothness and correspondence certainty. Instead we use coherent surfaces as more robust features to match.

### III. REGISTRATION ALGORITHM

The outline of our framework is shown in Figure 1. Different pipelines are used for the pre- and post-contrast images. The post-contrast image is used to build a patient-specific 3D shape model of the kidney using an off-the-shelf segmentation algorithm (Sec. III-A). On the other hand, the pre-contrast image is processed to extract edgel surfaces (Sec. III-B) which are then matched against the

patient-specific 3D shape model. The surface segmentation algorithm is one of the contributions of our work; it uses iteratively redefined curvature constraints to group cohesive sets of 3D edgels. A registration refinement stage is applied last. The key idea to our framework is the efficient combination of a bottom-up hypothesis generation mechanism (surface segmentation), and a top-down evaluation procedure (surface matching). Each of these steps are described next.

#### A. Model Segmentation (Post-contrast image)

Image segmentation is an effective way to build a detailed patient-specific shape model, although the associated difficulties often make this procedure unfeasible in practice. However, because of the inherent contrast enhancement in one of the images of our datasets, kidneys have enough saliency for reliable segmentation.

The input to our segmentation algorithm, besides the contrast enhanced image of course, is a rough bounding box containing the kidney and a rough scribble on kidney tissue in one slice. These two features can be moused in by an operator within a few seconds. The output is a binary 3D image where each voxel has a kidney or non-kidney label.

Segmentation is achieved using the Graph-Cuts [11] framework, in which a 3D image is represented with a graph. Each voxel becomes a node in the graph; we call these voxel-nodes. In addition, two synthetic nodes represent the foreground (*i.e.*, kidney) and background (*i.e.*, non-kidney) classes, which are referred to as class-nodes. The connectivity of the graph is as follows: each voxel-node is connected to its neighbors using a 6-connected pattern. Besides, each voxel-node is also connected to both class-

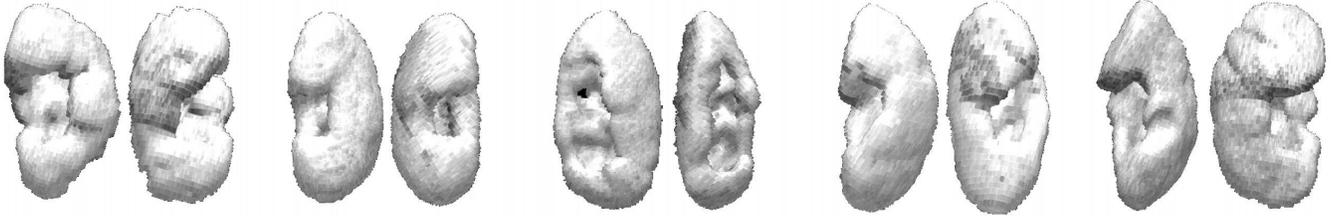


Figure 2: Segmentations of five kidneys from our dataset. Two views each.

nodes using two more edges, for a total of eight edges per voxel-node.

While this graph architecture is standard for all images, the information that drives the segmentation procedure is stored in the edge potentials. Segmentation is achieved by cutting the subset of edges with the smallest combined potentials (min-cut) to separate the foreground and the background nodes and their corresponding sub-graphs [12]. This formulation provides some intuition on what the properties of the edge potentials should be; edges that connect neighboring voxel-nodes ( $n_i$  and  $n_j$ ) should have potentials proportional to the similarity between the image intensities ( $I(x)$ ) of their corresponding voxels ( $x_i$  and  $x_j$ ). In our implementation, this function is

$$q(n_i, n_j) = \exp\left(-\frac{(I(x_i) - I(x_j))^2}{2\sigma^2}\right), \quad (1)$$

where  $\sigma$  determines how quickly the strength of the potential decreases as the intensity difference increases (in all our experiments we used  $\sigma = 10$ ).

On the other hand, edges that connect voxel-nodes with the foreground/background class-nodes should have potentials proportional to the likelihood of being part of the foreground/background classes. In our case, these potentials are computed using gray-level histogram distributions from data samples of each class. The foreground class sample is built from the moused-in scribble and its immediate 3D neighbors. The background class sample comes from the voxels on the moused-in bounding box of the kidney.

The method has been successfully used in this and many other settings, although the technical details are out of the scope of this paper. Figure 2 shows two views from five segmentations using our potentials and an off-the-shelf implementation [13]. Our experimental results showed that correct segmentations were obtained for all of our contrast enhanced images. Note that, due to the presence of lesions, the kidneys have a wide variety of shapes.

### B. Surface segmentation (Pre-contrast Image)

Unlike with the contrast enhanced images, segmentation of the entire kidney from the pre-contrast image is not feasible in practice, due to the common presence of regions with little or no contrast between the kidney and the surrounding

tissues. Instead, we propose computing segmentations of some salient surfaces, which despite their partial coverage, we hypothesized are sufficient for alignment. This hypothesis was validated by our experimental results.

Our surface segmentation algorithm is initialized by grouping cohesive sets of edgels from a small surface neighborhood via connected components. The algorithm then iterates between two steps: growing the patch according to some shape constraints, and updating these constraints. The steps are illustrated in Figure 3, using a real-data example, and are explained in detail below. After each run, the output is one segmented surface. The algorithm is repeatedly executed to extract more surfaces. We found that using a rough bounding box that contains the kidney is beneficial to decrease the number of non-kidney surfaces proposed by the algorithm, as well as the total computation time.

1) *Initialization*: A seed edgel is randomly chosen (Fig. 3a). Initially, seeds are sampled from a uniform distribution over all edgels, and it is possible that the selected seed edgel may not lie on the kidney surface. Nonetheless, the seed is used to grow a (kidney or non-kidney) surface as follows. Using a 6-neighborhood, 3D connectivity pattern, all edgels within a small radius of the seed are grouped, forming the initial surface patch  $P$ . A sphere is least-squares fit to this patch, optimizing for the radius and the position of the origin (Fig 3b). This sphere defines the initial shape model  $Q$ , which at this stage only captures a rough notion of the curvature of  $P$ , but that in subsequent iterations will become a more complex shape model.

2) *Constrained growth*: Using the same connectivity pattern as in the previous step, we build a list  $L$  of edgels that are neighbors to the currently segmented surface  $P$ . These become the candidates for inclusion to  $P$  at the current iteration. Each candidate is added to  $P$  only if closely extrapolated by the surface model  $Q$ . After evaluating all the candidates, the model  $Q$  is refit to the updated set of edgels in  $P$  (Figures 3c, d, details below). The process of recruit and refit is iterated. The algorithm stops when  $L$  is empty, or when none of the voxels in  $L$  are closely extrapolated by  $Q$ . At that point the current patch  $P$  is returned as the final segmented surface (Fig. 3e, f).

Multiple surfaces are segmented by repeating this process

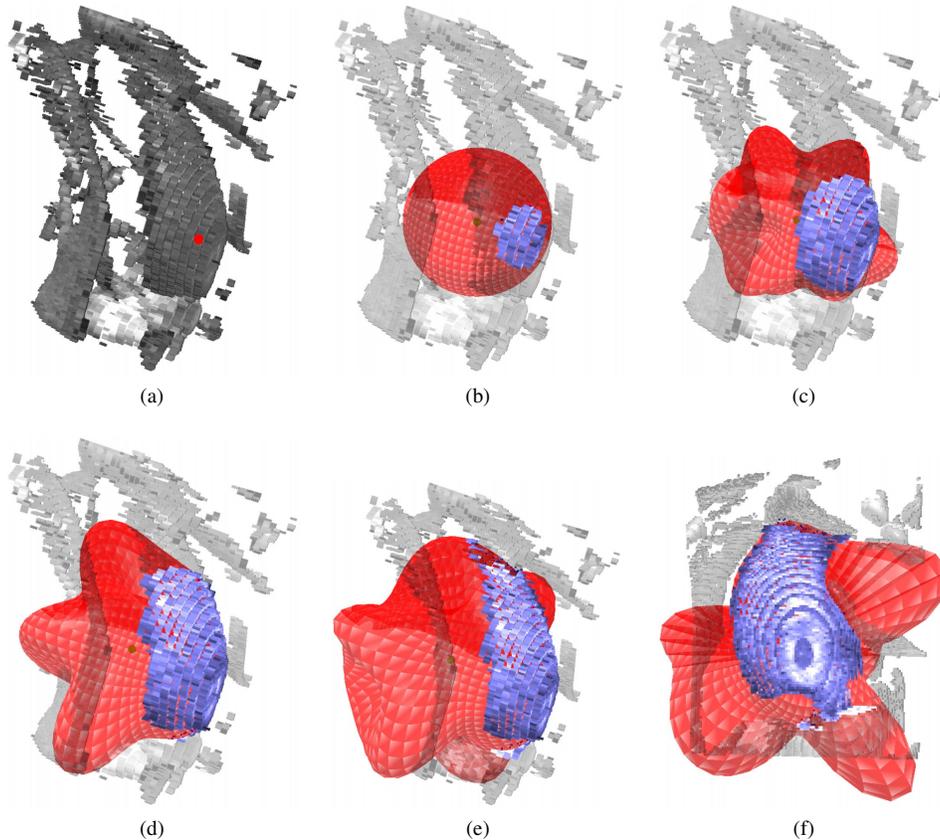


Figure 3: Surface segmentation: (a) Seed, in red; edgels in gray. (b) Initialization, model in red; initial patch in blue. (c-f) Surface segmentation iterations, SH model in red; segmented surface in blue. (f) Rotated view of (e). Note that the SH model is not expected to extrapolate to the entire kidney, but only to the neighboring voxels of the already fitted surface (in blue), see text.

of sampling a seed edgel and growing a surface patch. To encourage different surfaces to be extracted, we decrease the likelihood of sampling seeds from edgels which lie on previously extracted surfaces. The fitting algorithm, as well as the type of model used in the growth of these segmented surfaces are described next.

3) *Patch modeling*: Immediately after initialization, the patch  $P$  is small by construction, and can usually be well approximated by a spherical surface, but it grows after each iteration and will eventually describe a more complex geometry. The model  $Q$  must then be able to accommodate shape variability within a certain level of detail. We found that a Spherical Harmonics (SH) parameterized surface is a reasonable modeling technique, as one can adjust the level of detail by limiting the number of coefficients, much like in Fourier analysis. Spherical Harmonics are a smooth, single-valued family of functions. They can be used to describe star-shaped [14] or simply connected 3D [15] objects. In recent years, the computer vision community has used them to perform various kinds of shape analysis and modeling

[16], [17], [18].

While a SH model defines a surface over the entire domain of the sphere, it is important to notice that we do not expect to use the SH model to extrapolate the surface well away from the model’s support region. In fact, the example in Figure 3 shows that, because the support region spans only a local subset of the spherical angular domain, regions that are distant from the support region are very loosely constrained. However, because a SH model defines a smooth function, its shape should extrapolate to voxels adjacent to the patch, under the assumption that kidney surfaces are also smooth. We are only using the SH model to approximate the currently segmented patch and to extrapolate to its immediate neighboring voxels.

4) *SH model fitting*: Given a set of 3D points  $\{\mathbf{x}_i\}_{i=1}^N$ , a SH model can be extracted as follows. First, the origin of a polar coordinate system is placed at the center of the sphere that provides the best least squares fit. Then, each  $\mathbf{x}_i$  is represented in polar coordinates as  $\mathbf{s}_i = (\theta_i, \varphi_i, r_i)$ , with  $\theta_i \in [0, 2\pi)$  and  $\varphi_i \in [-\pi, \pi]$ . These points are then

approximated by a SH expansion of the form

$$r_i \approx \sum_{\ell=0}^L \sum_{m=-\ell}^{\ell} c_{\ell}^m Y_{\ell}^m(\theta_i, \varphi_i), \text{ for } i = 1, \dots, N, \quad (2)$$

where  $c_{\ell}^m$  are coefficients weighting the indexed spherical harmonic functions  $Y_{\ell}^m$ . In matrix notation, Equation 2 can be written as  $\mathbf{r} \approx \mathbf{Y}\mathbf{c}$ , with  $\mathbf{r} = (r_1, r_2, \dots, r_N)^T$ ,  $\mathbf{c}$  an  $M$ -vector consisting of the coefficients  $c_{\ell_j}^{m_j}$ ,  $j = 1, \dots, M$ , and  $\mathbf{Y}$  a  $N \times M$  matrix with the  $i, j$  entry is given by  $Y_{\ell_j}^{m_j}(\theta_i, \varphi_i)$ .

Typically the number of data points  $N$  is much larger than the the number of coefficients  $M$ , and therefore a least squares approach is used to solve for  $\mathbf{c}$ . The regularized objective function is:

$$\Phi(\mathbf{c}) = \|\mathbf{Y}\mathbf{c} - \mathbf{r}\|^2 + \mathbf{c}^T \mathbf{B}\mathbf{c}. \quad (3)$$

The first term corresponds to the reconstruction error. The second term is the regularizer that avoids overfitting by constraining the amount of deformation of the reconstructed surface. The matrix  $\mathbf{B}$  is diagonal. The value of  $b_{11}$  is always zero since  $c_1$  is related to a general surface scaling factor. Entries  $b_{ii} = \lambda$  (with  $i = \{1, 2, \dots, M\}$ ) control the penalty associated to  $c_i$  as they deviate from zero. The coefficients  $\mathbf{c}$  that minimize (3) are given by

$$\mathbf{c} = (\mathbf{Y}^T \mathbf{Y} + \mathbf{B})^{-1} \mathbf{Y}^T \mathbf{r}. \quad (4)$$

### C. Surface classification

At this point we have a set of surface patches that have been generated by the surface grouping algorithm described above, along with the 3D kidney model obtained from the post-contrast image. In order to determine whether any of these surfaces belongs to the kidney, all surfaces are putatively matched against the model. The problem is approached by first finding the transformation that maximizes the overlap between the model and the candidate surface, and then evaluating how much of an overlap there is.

1) *Model-surface fitting*: Let  $\mathcal{S}$  be a candidate surface with  $K$  edgels located at  $S = \{\mathbf{s}_k\}$ . The rigid transformation parameters ( $[\hat{R}|\hat{\mathbf{t}}]$ ) that brings this surface into alignment with the model are found using

$$[\hat{R}|\hat{\mathbf{t}}] = \operatorname{argmin}_{[R|\mathbf{t}]} \sum_{\mathbf{s} \in S} \rho(D(R\mathbf{s} + \mathbf{t})), \quad (5)$$

where  $D(\mathbf{x})$  is (an approximation to) the shortest Euclidean distance between  $\mathbf{x}$  and the patient-specific 3D model, and  $\rho(\cdot)$  is the Geman-McClure robust estimator [19], used to reduce the effect of outliers by limiting the maximum penalty. A quasi-Newton optimization method is used to find the (local) optimum set of parameters. To increase the likelihood of convergence to the function's global optimum, the optimization procedure is repeatedly run with varying initial guesses of the transformation parameters.

2) *Overlap metric*: Corresponding surfaces should coherently overlap with the 3D shape model. To asses this overlap, the candidate surface  $\mathcal{S}$  is first aligned onto the model using the parameters  $[\hat{R}|\hat{\mathbf{t}}]$  obtained from Equation 5 and then its edgels are evaluated using a function that rewards proximity and penalizes misalignment, namely

$$f(\mathbf{s}, [\hat{R}|\hat{\mathbf{t}}]) = \left[1 - \rho\left(D(\hat{R}\mathbf{s} + \hat{\mathbf{t}})\right)\right] \left(\hat{R}\mathbf{n}_{\mathbf{s}} \cdot \mathbf{n}_{m_s}\right). \quad (6)$$

The first term monotonically decreases with increasing distance from the surface edgel to the model. The functions  $D(\cdot)$  and  $\rho(\cdot)$  are the same distance and robust estimator functions, respectively, defined above. The second term penalizes misalignment between the edgel's normal ( $\mathbf{n}_{\mathbf{s}}$ ) and the model's normal at the closest model point ( $\mathbf{n}_{m_s}$ ). The cumulative evidence in support of the candidate surface as being part of the model is

$$\mathcal{O}(\mathcal{S}) = \sum_{\mathbf{s} \in \mathcal{S}} f(\mathbf{s}, [\hat{R}|\hat{\mathbf{t}}]). \quad (7)$$

The surface with the largest overlap score ( $\mathcal{O}(\mathcal{S})$ ) provides a good approximation to the transformation parameters ( $[\hat{R}|\hat{\mathbf{t}}]$ ). However, due to potential surface oversegmentation,  $\mathcal{S}$  may not be the only available kidney surface. The refinement stage (below) improves the transformation parameters by incorporating kidney edgels not present in  $\mathcal{S}$ .

### D. Refinement

In this stage the pre-contrast image edgels are first aligned to the kidney shape model using the estimated  $[\hat{R}|\hat{\mathbf{t}}]$ , and then the registration parameters are improved using those edgels that are sufficiently proximal and in good alignment with the shape model.

These two properties are evaluated using Equation 6. Edgels for which  $f(\mathbf{s}, [\hat{R}|\hat{\mathbf{t}}]) > k$ , are considered suitable and define the final support set  $\hat{\mathcal{S}}$ . We used  $k = 0.5$  for all our experiments.

The final registration transformation parameters are computed using

$$[R^*|\mathbf{t}^*] = \operatorname{argmin}_{[R|\mathbf{t}]} \sum_{\mathbf{s} \in \hat{\mathcal{S}}} \rho(D(R\mathbf{s} + \mathbf{t})). \quad (8)$$

The same optimization technique used to minimize Equation 5 was used here.

The average running time of the entire pipeline is kept under one minute on a conventional desktop computer; details are provided below.

## IV. EXPERIMENTAL RESULTS

To test the proposed approach, a set of 20 studies from real patients with some type of kidney lesion was used. To increase the efficiency of the images' dynamic range, we followed the standard procedure of scaling volumetric attenuation in (Hounsfield units HU) with a truncated linear function to compute image intensities; a window level of

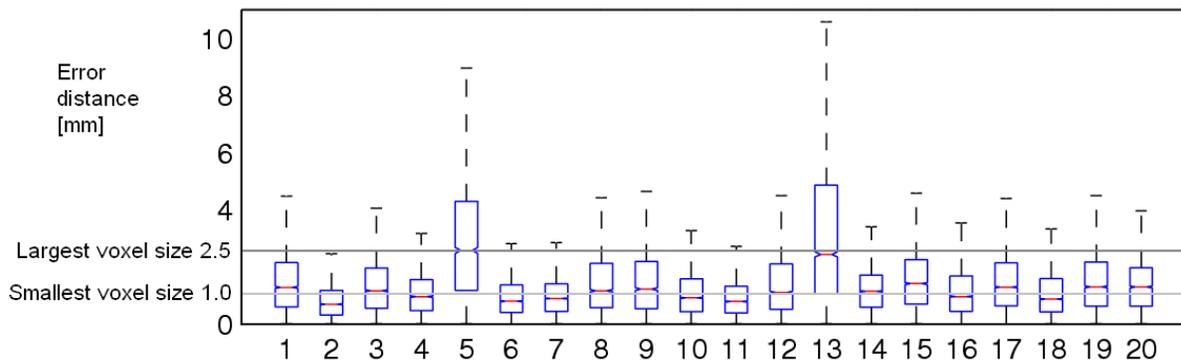


Figure 4: Box-plot of the distribution of distances between matched edges. The results of 18 out of the 20 test cases have more than three quarters of the validation surface samples closer to the segmented model than the largest voxel size

40HU and a window width of 400HU were used. Post-contrast images were also processed with bilateral filtering [20] to smooth consistent regions while preserving edges. The average size of a rough bounding box that contains a kidney is around  $120 \times 120 \times 60$  voxels, with a voxel size of  $1 \times 1 \times 2.5$  mm.

Experiments were run using a non-optimized implementation running on a Pentium IV desktop at 3.8 GHz with 2 GB of RAM and we found that the average running time could be kept under one minute. This time includes executing all the necessary image pre-processing steps (10-20 seconds), computing the 3D patient-specific model from the contrast enhanced image (2-5 seconds), running the surface segmentation module on the pre-contrast image several times (under 1 second per surface), and finding the final alignment (1-3 seconds).

The manual input required by the algorithm was provided by a user with no medical training other than the ability to identify the renal anatomy on a CT scan. This input consists of rough bounding boxes containing the kidney in both pre- and post-contrast images, as well as a scribble on the kidney in a single slice of the post-contrast image.

Edges are computed using a standard Canny edge detection algorithm, which uses 3D derivative-of-Gaussian kernels to compute gradients, followed by non-maximum suppression.

For the surface segmentation algorithm, the number of coefficients used in the spherical harmonics model implicitly determines the level of detail of the reconstructed surface. We empirically found that the algorithm performed best with  $m = 25$  coefficients (Equation 2) regularized by  $\lambda = 0.025$  (Equations 3 and 4) for this type of data. This combination of parameters exhibits good extrapolation properties on the smoother regions (of both kidneys and the surrounding structures) while preventing the surface segmentation procedure from under segmenting. The quasi-Newton BFGS method

was used to fit the sphere parameters (*i.e.*, the center and radius) for the surface segmentation algorithm. The same method was used to optimize the parameters of the alignment between candidate surfaces and the 3D model.

1) *Quantitative validation:* The whole registration pipeline was applied to each of the 20 image pairs in our dataset. To quantitatively evaluate the quality of the resulting transformation, all the detected edges from the pre-contrast image were evaluated using Equation 6. Edgels with  $O(\mathbf{s}_k, [\hat{R}|\hat{\mathbf{t}}]) > k_v$ , were labelled as kidney surface inliers, and defined the validation set  $T$ . The threshold  $k_v$  was even more tightly defined ( $k_v = 0.9$ ) in order to more clearly assess the alignment. The distance from an edgel in  $T$  (from the pre-contrast image) to the closest edgel from the segmented model (from the post-contrast image) was computed. Ideally, a perfectly registered image pair should produce zero distances between corresponding edges. Figure 4 shows a box plot of edge-to-model distances, where 18 out of the 20 test cases have three quarters of the validation surface samples closer to the segmented model than the (largest) voxel size.

The same figure shows two failure cases; the first one due to an abnormal structure of size comparable to the kidney itself, which altered most of the pre-contrast image edge information and led to a poor model fit. The second one presents an extremely narrow kidney, inaccurately modeled using the current level of detail.

2) *Qualitative validation:* Aligned sets were given to a specialized radiologist (a co-author), who provided feedback on the accuracy of the registered sets. According to the expert, the method provides the means to much more rapidly identify and locate higher risk regions. Two series of the resulting registered images are shown in Fig. 5. The leftmost column corresponds to the pre-contrast images, and columns to the right show post-contrast images progressively more strongly overlaid. For reference, contours of the kidneys

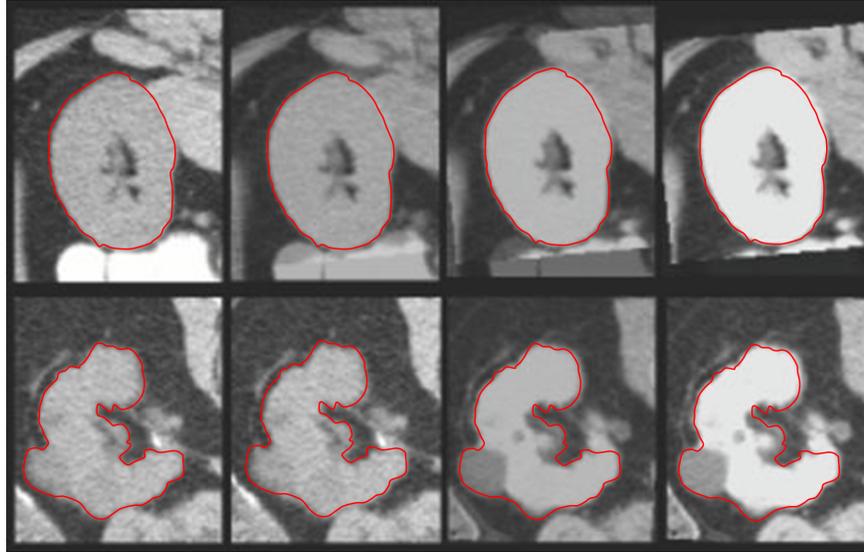


Figure 5: Qualitative validation: overlaid alpha-blended registered images. The leftmost column corresponds to the pre-contrast images, and columns to the right show post-contrast images progressively more strongly overlaid. For reference, contours of the kidneys from the low-contrast images are displayed on all overlays.

from the low-contrast images are displayed on all overlays.

#### V. SUMMARY

We present a method to register kidneys from volumetric images (CT scans) by finding the alignment between automatically segmented surface patches and a 3D model. We also present an algorithm to automatically segment coherent and smooth surfaces from edgel maps. Our non-optimized implementation runs within seconds on a single CPU, and requires minimal user intervention. Validation was performed using image pairs from real patients with different degrees of kidney damage. Quantitative validation showed that most registration errors are smaller than the largest voxel size for 90% of the test cases. In the clinical domain, the results were qualitatively validated by an expert radiologist as useful for diagnosis.

#### REFERENCES

- [1] J. Maintz and M. Viergever, "A survey of medical image registration," *Medical Image Analysis*, vol. 2, no. 1, pp. 1–36, 1998.
- [2] C. J. Twining, T. Cootes, S. Marsland, and C. J. Taylor, "A unified information-theoretic approach to the correspondence problem in image registration," in *ICPR*, 2004, pp. 704–709.
- [3] S. M. Pizer, G. Gerig, S. Joshi, and S. R. Aylward, "Multiscale medial shape-based analysis of image objects," *Proceedings of the IEEE*, vol. 91, no. 10, pp. 1670–1679, October 2003.
- [4] P. J. Besl and N. D. McKay, "A method for registration of 3-D shapes," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 14, pp. 239–256, February 1992.
- [5] Y. Boykov, V. S. Lee, H. Rusinek, and R. Bansal, "Segmentation of dynamic N-D data sets via graph cuts using markov models," in *In Proc. Medical Image Computing and ComputerAssisted Intervention*, 2001, pp. 1058–1066.
- [6] P. A. Viola and W. M. W. III, "Alignment by maximization of mutual information ." in *ICCV'95*, 1995, pp. 16–23.
- [7] Y. Sun, M.-P. Jolly, and J. M. F. Moura, "Contrast invariant registration of cardiac and renal MR perfusion images," *MICCAI*, vol. 1, pp. 903–910, 2004.
- [8] H. Huang, L. Shen, R. Zhang, F. Makedon, B. Hettleman, and J. Pearlman, "Surface alignment of 3D spherical harmonic models: Application to cardiac MRI analysis," in *MICCAI*, 2005, pp. 67–74.
- [9] S. Heldmann and S. Zidowitz, "Elastic registration of multiphase CT images of liver," in *SPIE Medical Imaging 2009: Image Processing*, vol. 7259. SPIE, March 2009.
- [10] D. Kwon, I. D. Yun, K. H. Lee, and S. U. Lee, "Efficient feature-based nonrigid registration of multiphase liver CT volumes," in *In BMVC*, 2008.
- [11] Y. Boykov and G. Funka-Lea, "Graph cuts and efficient N-D image segmentation," *IJCV*, vol. 70, no. 2, pp. 109–131, 2006.
- [12] Y. Boykov and M. P. Jolly, "Interactive graph cuts for optimal boundary and region segmentation of objects in N-D images," *ICCV*, vol. 1, pp. 105–112, 2001.
- [13] <http://vision.csd.uwo.ca/code/>
- [14] W. Magnus, F. Oberhettinger, and R. Soni, *Formulas and Theorems for the Special Functions of Mathematical Physics*, 1966.

- [15] C. Brechbühler, G. Gerig, and O. Kübler, "Parametrization of closed surfaces for 3-D shape description," *Computer Vision and Image Understanding*, vol. 61, pp. 154–170, March 1995.
- [16] J.-L. Dillenseger, H. Guillaume, and J. J. Patard, "Spherical harmonics based intrasubject 3-D kidney modeling/registration technique applied on partial information," *IEEE Transactions on Biomedical Engineering*, vol. 53, no. 11, pp. 2185–2193, Nov. 2006.
- [17] M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz, "Rotation invariant spherical harmonic representation of 3D shape descriptors," in *Eurographics/ACM SIGGRAPH symposium on Geometry processing*, 2003, pp. 156–164.
- [18] R. J. Morris, R. J. Najmanovich, A. Kahraman, and J. M. Thornton, "Real spherical harmonic expansion coefficients as 3D shape descriptors for protein binding pocket and ligand comparisons," *Bioinformatics*, vol. 21, pp. 2347–2355, 2005.
- [19] M. J. Black and A. Rangarajan, "On the unification of line processes, outlier rejection, and robust statistics with applications in early vision," *IJCV*, vol. 19, no. 1, pp. 57–91, 1996.
- [20] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," in *ICCV*, 1998, p. 839.
- [21] E. Trucco and A. Fusiello and V. Roberto, "Robust motion and correspondence of noisy 3-D point sets with missing data," in *Pattern Recognition Letters*, vol. 20, no. 9, p. 889–898, 1999.
- [22] L. M. Lui, T. W. Wong, P. M. Thompson, T. F. Chan, X. Gu, and S.-T. Yau, "Shape-based diffeomorphic registration on hippocampal surfaces using beltrami holomorphic flow." in *MICCAI (2)'10*, 2010, pp. 323–330.