Deformable Model-Based Segmentation of 3D CT by Matching Distributions

A rate-limiting step in tomographic image guided therapy or 4D CT-based treatment planning is the significant amount of time and human intervention required to delineate the tumor and nearby critical structures on each scan. A radiation oncologist can often take 30-45 minutes to outline all of the structures of interest in every axial slice. Performing this operation each of the 30-40 times a patient is treated is tedious. Using an uncompiled MATLAB implementation on a modest hardware platform (1.67 GHz AMD machine with 448 MB RAM), we present an algorithm that performs the same procedure in less than five minutes. Our early results show in-plane agreement with manual segmentation of multiple organs in the male pelvis on the order of 2-3mm.

The state of the art in computer vision algorithms applied to prostate segmentation is considerably less advanced than for other sites. Manual contouring remains the gold standard. This is due to several algorithmic challenges inherent in segmenting deformable objects from real-world 3D images:

- Challenge #1: The objects of interest are often diffuse and lack strong edges.
- Challenge #2: There are often many objects, both of interest and not of interest, within a small volume.
- Challenge #3: Many objects have fairly similar intensity profiles. Typically, this effect cannot be removed by simple pre-processing such as histogram equalization.
- Challenge #4: Many of the objects are of roughly the same shape. For example, the prostate and bladder are both "somewhat deformed" spheres.

The algorithm we present uses learned models for both the shape and appearance of objects to achieve segmentation; learning both types of information is the only reasonable way to deal with all 4 challenges. Existing algorithms that use both shape and appearance models (such as [Cootes01]) require a pixelwise correspondence between the model and the image; this correspondence is often very difficult to compute, and can be extremely time-consuming. Instead, our algorithm characterizes a model object by (a) its shape and (b) a probability distribution of the intensities (or colors, textures) of the pixels within its interior. As a result, comparing a particular model object to the image is as simple as comparing two probability distributions. The basic idea is as follows. The shape is given by a description of the surface, or multiple surfaces in the case of multi-object segmentation. The appearance is described by a probability distribution of some photometric variable inside the object of interest, or multiple distributions in the case of multi-object segmentation. A shape-appearance pair is then given by (surface, distribution), and this pair is considered sufficient to characterize an object for the sake of segmentation. The learned model is a low-dimensional manifold in the space of such pairs. To verify how well any particular shape-appearance pair matches the image, we compute the empirical distribution of the photometric variable inside the shape within the image; this distribution is then compared to the appearance model. We therefore evolve the shape of an object (or multiple objects) until the empirical distribution(s) best matches the model distribution(s), resulting in a final segmentation.

Our training set consisted of 21 sets of 512x512x90 CT images of the male pelvis from different patients (under bladder-full conditions). In each of the images, a radiation oncologist had outlined the prostate, bladder, and anterior rectal wall. There is a substantial amount of inter-patient shape variability in this data set. We resampled each object to have 20 slices with 20 points on each contour and built two PCA shape models: one model using the prostate alone, and one joint model for the bladder, prostate, and anterior rectal wall, since the bladder and rectum are radiation-sensitive structures that should be detected and avoided during radiation delivery. Figure 1 shows the effect of varying the two most dominant modes of the joint model. We can see that these models are able to capture nonrigid size and shape variability. We initialized the model at the position of the mean shape and allowed it to converge using the distribution matching algorithm. Note that this is in contrast to other algorithms where rough manual placement is required to guarantee convergence to the correct result. The results for the joint shape model applied to one of the 21 patients are shown in Figures 2 and 3.

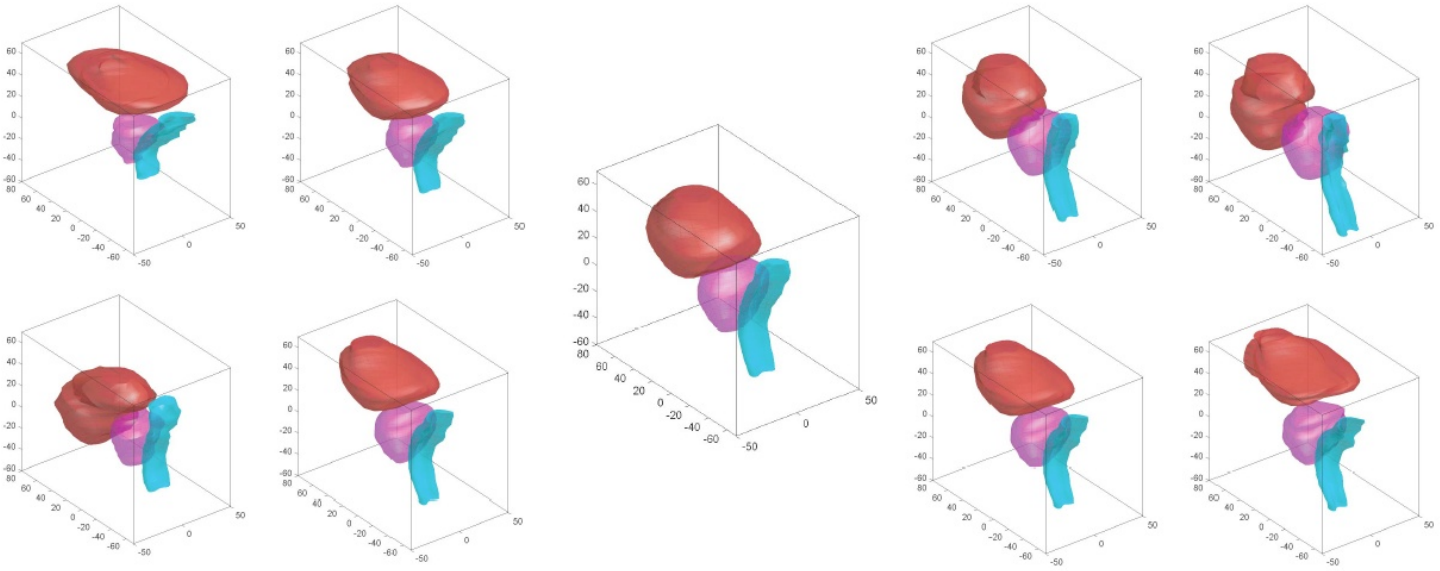

Center: The mean shape of the PCA model of the bladder, prostate, and rectum. The rows correspond to varying the two most dominant modes of the joint shape model of the bladder, prostate, and anterior rectal wall.

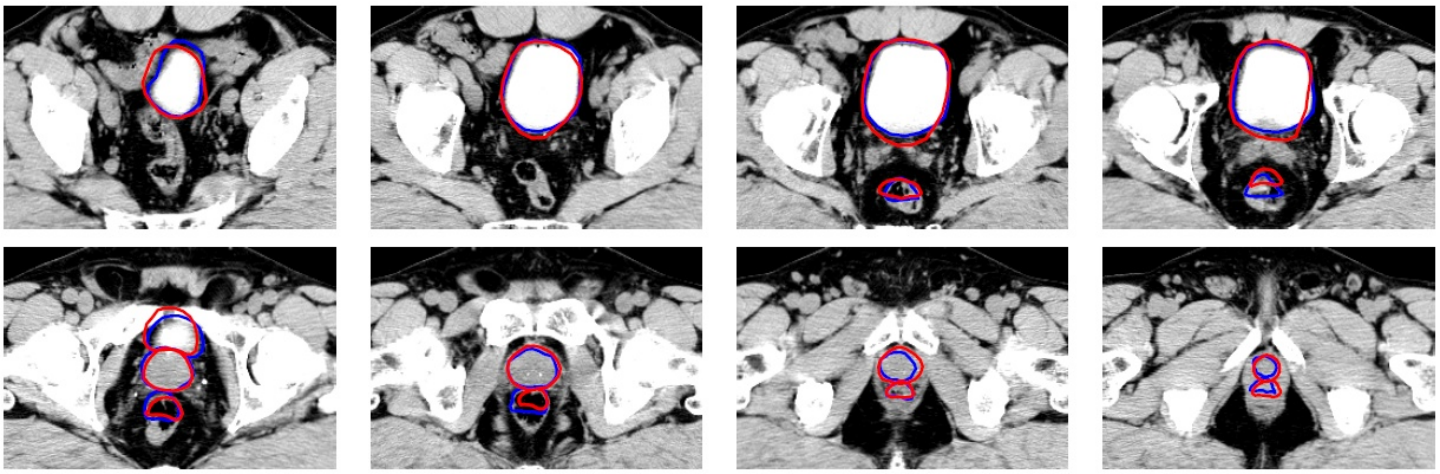

Segmentation results for evenly-spaced slices of patient 2664 using the joint object model (10 modes). The red contour shows the segmentation result at convergence. The blue contour shows the hand-drawn ground-truth contours supplied by a radiation oncologist. The segmented organs are, from top to bottom, the bladder, prostate, and anterior rectal wall. Note that each organ is not visible in every slice.

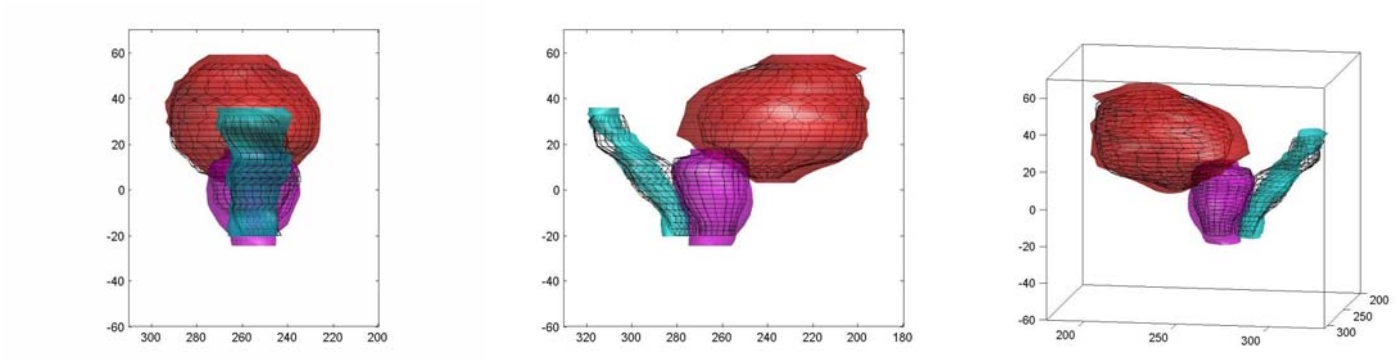

(a) AP, (b) lateral and (c) oblique 3D visualizations of the joint model (solid) vs. ground truth (wireframe).