

CHAPTER 13: EFFICIENT KERNEL DENSITY ESTIMATION OF SHAPE AND INTENSITY PRIORS




Figure 1. Segmentation challenges and estimated intensity distributions. The two curves on the right correspond to the empirical probability of intensities inside and outside the left ventricle (for the ultrasound image) and the prostate (for the CT image). The region-based segmentation of these structures is a challenging problem, because objects and background have similar histograms. Our segmentation scheme optimally exploits the estimated probabilistic intensity models.




Figure 2. Schematic plots of different density estimates within a subspace. Darker shading indicates areas of high probability density for the respective models. The kernel density estimator adapts to the training data more flexibly since it does not rely on specific assumptions about the shape of the distribution.




Figure 3. Linear versus nonlinear shape interpolation. The upper row shows 6 out of 49 training shapes and a 3D projection of the isosurface of the estimated (48-dimensional) shape distribution. The latter is clearly neither uniform nor Gaussian. The bottom row shows a morphing between two sample shapes along geodesics induced by a uniform or a kernel distribution. The uniform distribution induces a morphing where legs disappear and reappear and where the arm motion is not captured. The nonlinear sampling provides more realistic intermediate shapes. We chose human silhouettes because they exhibit more pronounced shape variability than most medical structures we analyzed.




Figure 4. Segmentation without prior. Since there is no shape constraint imposed upon the contour — other than a small length constraint present in the Chan-Vese model — the boundary gradually separates brighter from darker areas. This indicates that intensity information is insufficient to induce the desired segmentation.

CHAPTER 13: EFFICIENT KERNEL DENSITY ESTIMATION OF SHAPE AND INTENSITY PRIORS




Figure 5. Boundary evolution for an ultrasound segmentation with uniform shape prior. By constraining the level set evolution to the linear subspace spanned by the first few eigenmodes computed from a set of training shapes, one can improve the segmentation of the given image (see, e.g., [?]). Nevertheless, in our application, the uniform shape prior does not sufficiently constrain the segmentation process, permitting the boundary to leak into darker image areas.




Figure 6. Boundary evolution for an ultrasound segmentation with nonparametric shape prior. Imposing a non-parametric shape prior within the eigenspace spanned by the training shapes leads to a segmentation process that is sufficiently constrained to enable an accurate segmentation of the left ventricle. In contrast to the uniform prior (see Figure 5, the nonparametric one does suppress leaking of the boundary), because it constrains the level set function to a well-defined submanifold around the training shapes (see also Figure 2).




Figure 7. Prostate segmentation for two patients with the same shape model. Each row shows axial slices of the same segmentation for one patient. The manual segmentation is in black and the automatic one white.




Figure 8. Comparison of the segmentations obtained with the kernel prior (white) and with alternative approaches (black).




Figure 9. Robustness to noise.