

Deep Structure of Images in Populations via Geometric Models in Populations

Stephen M. Pizer, Ja-Yeon Jeong, Robert E. Broadhurst, Sean Ho, Joshua Stough
Medical Image Display & Analysis Group
University of North Carolina, Chapel Hill, NC, USA

Abstract. We face the question of how to produce a scale space of image intensities relative to a scale space of objects or other characteristic image regions filling up the image space, when both images and objects are understood to come from a population. We argue for a schema combining a multi-scale image representation with a multi-scale representation of objects or regions. The objects or regions at one scale level are produced using soft-edged apertures, which are subdivided into sub-regions. The intensities in the regions are represented using histograms. Relevant probabilities of region shape and inter-relations between region geometry and of histograms are described, and the means is given of inter-relating the intensity probabilities and geometric probabilities by producing the probabilities of intensities conditioned on geometry.

1. Introduction

This workshop focuses largely on finding the essential structure in the 2D images that are input to the human visual system. It marvels that the mature visual system can in 100-200ms make out a complex object, a multi-object complex, or a textured region such as a forest and identify it, and it wonders how such direct image access is possible.

But the immature human visual system has a far more limited capability, and some part of that limitation is that it has not built up the models of the world corresponding to the objects and regions of texture that can be recognized. The models are somehow a very important part of the deep structure. This position applies as well when we are exploring the methods of computer vision, or more generally, of image analysis. This paper therefore explores a model-relative point of view for which the deep structure is found in an image when such models are present.

First, where do the models come from? They must come from multiple instances of the geometric layouts, i.e., using statistics to form probabilistic descriptions. Similarly, the images must come from multiple instances. We therefore expect that the structure we build will involve *probability densities on geometry and probability densities on images*. The model-relative point of view suggests that we will need probabilities on images \underline{I} *conditioned on geometry* \underline{Z} .

For reasons that need not be repeated to attendees of this workshop, the image data need to be understood at multiple scales. However, we argue that this data need to be

understood in reference to geometric models that themselves are understood at multiple scales. It is this interaction between image intensity data and geometric models that is the essence of this paper. We do not claim to have solved the problem completely, but we hope to communicate some useful points and to develop some important relations. We have developed these largely in the context of our main application of the analysis of 3D medical images such as CT and MR images, but we suggest that many of the ideas apply in the context of 2D light images of the visible world (e.g., see the work on active appearance models by Cootes et al. [1998]).

Sections 2-4 describe the basic ideas of the geometric descriptions and relations, intensity descriptions, and the probabilistic relations between geometry and intensity that we propose. Section 5 presents the ideas of geometry and its probability with mathematical detail, and section 6 presents the ideas of image intensities conditioned on geometry and its probability. Section 7 describes evidence for the usefulness of the proposed representation, and section 8 closes this paper with a discussion of issues yet to be faced and work yet to be done.

2. Objects, object-based coordinates, and neighbors

In images, objects or named real world regions are typically characterized by a uniform color or color mixture or texture that may vary in a predictable way across the object, or they are made up of a small collection of such regions. Thus the regions are formed, to first approximation, by successive subdivision. A major point is that one needs a coordinate system *relative to the object or region* to describe this intensity information in the context of a collection of objects or real world regions that vary in shape across the population of such objects. To save space, henceforth we will call this an “object-based coordinate system”. With such a coordinate system \mathbf{u} , intensity properties \mathbf{I} can be transformed from the Euclidean coordinates \mathbf{x} in which they arrive to the camera to the object-based coordinates, producing an intensity distribution $I(\mathbf{u})$ that can be the subject of statistical analysis. Similarly, objects, named regions, *and their inter-relations* have a statistical geometry that is coupled with the image intensities or textures that yield the object percept.

Inter-relations among neighbors are central to geometry and thus images of geometric entities. But these relations occur at many scales: objects and named regions have neighbors that are other objects and named regions; object sections have neighbors that are other object sections; voxels or pixels have neighbors that are other voxels or pixels. This setup is illustrated in Fig. 1.

If describing the inter-relations among the smaller scale entities is not to produce a combinatorically unwieldy representation, these neighborhood relations need to be pretty local *relative to their scale*. Thus, we need multiscale geometry and multiscale image intensity representations, and we need a geometry that understands abutments and region inter-relations at a scale.

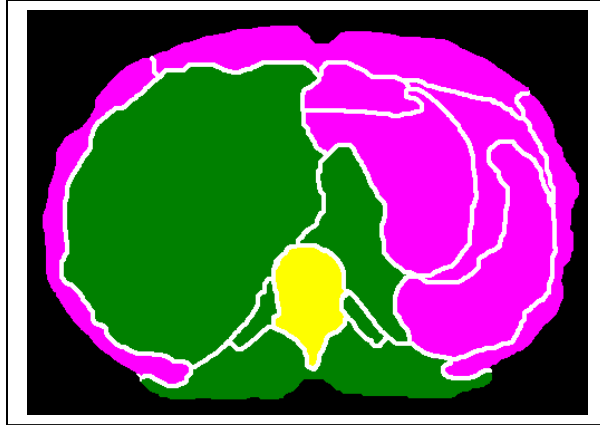


Fig. 1. An abdominal image region made up of sub-regions, showing only one level of subdivision. A region of interest, shown in the lightest shade, can be the whole image domain or a subset of that domain that itself forms a region at some scale. The neighbors of the region of interest are shown in the darkest shade.

For the purposes of this paper we will assume a fixed subdivision topology and a fixed neighborhood relation topology. Variation in this topology is a worthwhile topic but not within the realm of this paper. Maintaining this topology allows slidings of one region along another at the regional scale and these are needed, not just diffeomorphic transformations. Abutting objects or nearby objects can not only change relative shape, but also they can change relative position and orientation, in a way quite possibly correlated with the shape of their neighboring objects and regions. A way of probabilistically measuring a combination of object shape and inter-object geometric relation is discussed in the paper by Pizer, Jeong, Lu, et al. in this volume [2005].

3. Intensity histograms via object-based apertures

Regional intensity representations have, we believe, traditionally been far too local, indicating pixel by pixel what the probability distribution of intensity should be. Of course, summarizing intensity over not-too-local regions is the aim of using apertures in the aperture scale space methods of ter Haar Romeny [2003], Florack [2000], Lindeberg [1994], and many others. However, we argue first, that the scale selection should be with respect to the geometry-based coordinate system \mathbf{u} rather than the Euclidean coordinate system \mathbf{x} , and second, that the information within the aperture can usefully be more richly summarized as a histogram rather than as a single intensity for the aperture-weighted region. The first point is buttressed by results of Sean Ho's dissertation [2004] that probability distributions on intensities in an along-object-boundary scale space are more generalizable and specific than the more commonly used ones that are local in along-boundary location. The second point, regarding the benefits of the uses of probabilities on regional histograms, is discussed in the companion paper by Broadhurst et al. [2005].

That paper gives some early results on the method for computing statistics of intensity histograms and the advantages not only of that choice but of doing this across multiple regions chosen in object-relative coordinates. Of course, histogram probabilities are only useful if the image intensities are somehow normalized to remove intensity non-calibration effects, but this is true for any analysis of intensities across different images.

The regional, histogram-based approach, of course, begs the question of region size, and this is a matter of scale, i.e., aperture. Because the regions have tolerance at any scale, the aperture corresponding to a region needs to be soft with a falloff distance that increases with region size, but on the other hand, the aperture needs to be determined by the statistics of the intensity distribution in the region, and this may have a rather harder or rather softer edge. We discuss this further in section 4.

4. Geometric and histogram statistics

The geometric relations that we have been discussing, of the formation of object-based coordinates, are most economically understood in terms of local orientation at least as much as position. Indeed, the early human vision system has orientation, position, and scale as its three major coordinates. Together, these spaces, or the abstract spaces of the corresponding geometric transformations of rotation, translation, and magnification, are curved. Therefore, if we are to do probabilistic analysis in these terms, our Euclidean feature spaces will not do. Linear spaces may provide the best engineering mechanisms for statistics, and nonlinear transformations such as the rotations involved in local twisting and bending may have their mathematical underpinnings in tangent planes (linear spaces) on curved manifolds. Nevertheless, probabilities defined on the non-Euclidean feature spaces must be the essential basis for our geometric statistics, and these have been provided by the methods of Fletcher [2004] and Pennec [1999] that have been built on statistical approaches pioneered by Kendall [1989] and Grenander [1976, 1978]. Based on these ideas, we use probability distributions based on Fréchet means on curved manifolds and principal geodesics defined via manifold tangent planes at the mean.

Since we need a discrete scale space via successive subdivision, we need a means of producing probabilities at any scale. For this we use the method of residues described in [Lu 2005]. The idea is to describe the probability distribution for the global region first and define the residue at the global scale level for each training sample to be the representation of the training sample itself. Then successively, large scale to small (increasing j in Fig. 2), we remove from the residue of each training sample at a scale level j the projection onto the manifold forming the domain of the probability at scale level $j-1$. This leaves residues at each scale level that describe only information at that scale level, i.e., an analogy to the Laplacian pyramid. Probabilities can then be done scale level by scale level on the corresponding residues.

Histogram statistics need to be provided in a feature space in which straight lines (geodesics) between histograms produce valid interpolations. In the companion paper by Broadhurst et al. [2005], a histogram representation in terms of average quantile values of the histogram is described that has the desired property: moving along geodesics

whichever representation you might use, as long as that representation provides region-based coordinates.

Statistical analyses of intensities requires that intensities be compared at corresponding positions across training cases. Similarly, subdivision of regions must be done with appropriate spatial correspondence. These correspondences are implied by the geometry, i.e., by the layout and shape of the regions, perhaps via statistics of the training cases. We characterize this correspondence via the object-based coordinate system named \mathbf{u} .

Let \mathbf{I} describe a discrete image, i.e., a tuple made up of a scalar value for each pixel or voxel. The object-based coordinates of a region allow us to designate those pixels or voxels within the region and thus allow the formation of histograms of objects or other connected image regions. These objects or regions have a neighborhood relation with abutting regions. If histograms are to be the region descriptors, probably the smallest size regions should still have many pixels or voxels.

If the space described by \mathbf{z} is to be divided into subregions, each subregion can be described as an aperture, i.e., a weighting function in \mathbf{u} . We see the ideal aperture as representing the certainty a pixel or voxel is in the region. Thus the aperture would have the value of 1 well interior to the region “boundary”, have the value of 0 well exterior to the region “boundary”, and fall smoothly from 1 to 0. These spatial subdivisions forming the regions need to get successively smaller as the level of detail increases. The need for discrete subdivision suggests that the levels of detail be discrete and that the process is subdivision, i.e., decreasing aperture size or scale. Thus we index level by j , and we take the inverse point of view of many at this workshop and let j increase as the level of detail, i.e., the spatial scale decreases. Let \mathbf{z}_k^j indicate the aperture describing the k^{th} region at the j^{th} scale level. And let \mathbf{u}_k^j describe the object-based coordinate system implied by \mathbf{z}_k^j .

We need to describe the intensities within the aperture (region) \mathbf{z}_k^j . These can be described by histograms of $I(\mathbf{u}_k^j)$, with aperture weighted counts. We use the symbol H_k^j for the histogram for the k^{th} region at scale level j . Notice that, rather than being described directly in \mathbf{x} , \mathbf{z}_k^j needs to be described by apertures in \mathbf{u}^{j-1}_m , where m designates the parent region of region k,j . Each object-based coordinate must then be transferred to Euclidean coordinates \mathbf{x} .

Because regions not only have soft apertures but also have level surfaces of aperture weight that are smoother, the larger the scale (lower the scale index j), we can expect object regions to contain a small contribution from pixels or voxels in neighboring regions. Indeed, it is possible that we may wish to train histograms excluding these neighbor-region voxels and even possible that we might train boundary region histograms that explicitly have a nearly half-and-half mixture of the two neighboring histograms.

Just as regions have soft apertures in space, histograms need soft apertures in intensity. These apertures should increase as the spatial scale increases. In effect, this means that the bin sizes of the histograms should increase with spatial scale and that the bins should overlap, having soft edges. This might be done directly, but it is probably preferable to do

it on the Earth Mover's features of histograms produced by Gaussian-weighted average quantile values of histograms [Broadhurst 2005].

The regional geometric representations need to be described over a population of cases, and the regional intensity distributions need to be described over that population of cases, so we need probability densities on multiscale geometric histograms. The means of doing statistics on histograms is described in detail in the companion paper by Broadhurst et al. [2005].

As for statistics on geometry, besides understanding that statistics on geodesic manifolds will be necessary if twisting and bending is to be described and that these statistics will be on residues at the respective scales, it is important to understand that a region's geometry does not simply involve its own shape but also that of its neighbors. One of the major functions of having various scales is to allow neighbors to be local at that scale. Indeed, our picture is that a neighbor of a region includes all and only abutting regions at the same scale as the region. Thus denote by $N(j,k)$ the set of regions that are neighbors of region j,k : $\{j,m \mid \text{region } j,m \text{ abuts region } j,k\}$. Of course, this definition can be loosened to include second-order abutments or even higher order abutments. Since objects can slide along each other in these populations, we have another orientational relation that must be described using geodesic manifolds.

The importance of relationships with neighbors means that a Markov Random Field is the appropriate model of geometry. That is, we need to focus on $p(\mathbf{z}_k^j \mid \mathbf{z}_{N(j,k)})$, where we use the notation $\mathbf{z}_{N(j,k)}$ as a shorthand for $\{\mathbf{z}_m^j \mid (j,m) \in N(j,k)\}$. This density involves two factors, the shape of the region and its inter-relationship with its neighbors: $p(\mathbf{z}_k^j \mid \mathbf{z}_{N(j,k)}) = p(\text{shape of } \mathbf{z}_k^j, \text{inter-relationship of } \mathbf{z}_k^j \text{ and } \mathbf{z}_{N(j,k)})$.

This suggests handling object shape and inter-object relations together: one way to do this is via the region augmentation and prediction method described in the companion paper by Pizer, Lu, Jeong, et al. That paper shows how to train the shape plus neighbor probabilities. Also, [Han 2005] shows how to use this same approach for the relation between protrusion and indentation subfigures and their host figures. Rather than using these conditional probabilities to compute $p(\mathbf{z}^j)$, this latter probability density is never evaluated, but rather one can find the mode of $p(\mathbf{z}^j)$ by the Iterative Conditional Modes algorithm using the conditional probabilities $p(\mathbf{z}_k^j \mid \mathbf{z}_{N(j,k)})$.

As shown by the down-pointing arrow in the "Geometry" column in Fig. 2, we create the geometric statistics from large scale to small. In contrast, histograms are created from the original image, i.e., from the smallest scale, so the vertical arrow in the "Image Intensity" column is shown pointing upward. As described earlier, we use residues to create geometric probabilities for scale j based on removing the larger scale changes at scale $j-1$ from the training cases. We have thereby implicitly assumed the conditional independence of the probabilities of the geometry of a parent region with each of its subregions. But we have found empirically that this property does not hold. Thus we either must include parent regions and children regions as neighbors, or we must adjust

the principal geodesic analysis forming the probabilities to iteratively adjust the respective probabilities to obtain the desired property.

6. Mathematics and probability of coupled image histograms and multiscale geometry

We now move on to combining the intensity histogram statistics and the geometric statistics. As shown in Fig. 2, we believe it is appropriate to combine these scale level by scale level. Let us take segmentation, i.e., region designation in a particular target case, as our driving problem. From the Bayesian point of view we wish to compute the posterior optimum, i.e. $\arg \max_{\mathbf{z}} p(\mathbf{z} | \mathbb{I})$. Doing this scale by scale means optimizing $p(\mathbf{z}^j | \mathbb{I})$. Moreover, we do not optimize $p(\mathbf{z}^j | \mathbb{I})$ directly but rather apply the Iterative Conditional Modes algorithm to $p(\mathbf{z}_k^j | \mathbf{z}_{N(j,k)}, \mathbb{I})$. But $p(\mathbf{z}_k^j | \mathbf{z}_{N(j,k)}, \mathbb{I}) = p(\mathbb{I} | \mathbf{z}_{k \cup N(j,k)}^j) p(\mathbf{z}_k^j | \mathbf{z}_{N(j,k)}) \times$ a constant with respect to \mathbf{z}_k^j .

The density $p(\mathbb{I} | \mathbf{z}_{k \cup N(j,k)}^j)$ is related to the joint density $p(\mathbf{z}_{k \cup N(j,k)}^j, \mathbb{I})$ by $p(\mathbf{z}_{k \cup N(j,k)}^j, \mathbb{I}) = p(\mathbb{I} | \mathbf{z}_{k \cup N(j,k)}^j) p(\mathbf{z}_{k \cup N(j,k)}^j)$. The latter factor is a subset of the characterization of the geometry of the j,kth region and its inter-relation with its neighbors. Therefore, the intensity information, given the geometry can be considered to be fully independent of the geometry, once the intensity information is analyzed in object-relative coordinates. More precisely, this will be true except for the small effects of overlapping skirts of the apertures for adjacent regions. This proposition may well not hold for regions of just a few pixels or voxels, and for these the more common Gaussian aperture-based single intensity value might be the preferred basis for probabilities.

From the independence proposition it follows that $p(\mathbb{I} | \mathbf{z}_{k \cup N(j,k)}^j)$ can be given as a product of histograms relevant to scale j. Which histograms are these? Consider $j=1$, the scale level at which information global to the whole domain of the image is summarized. At that scale level, there is only one region, and histograms of each of its sub regions is needed, i.e., $p(\mathbb{I} | \mathbf{z}^1) = \prod_{i \in \text{regions at scale 2}} p(H_i^2)$.

A similar argument applies for smaller (larger index) scale levels. $p(\mathbb{I} | \mathbf{z}_{k \cup N(j,k)}^j)$ requires the histograms of all sub-regions of \mathbf{z}_k^j , H_m^{j+1} for m a child of j,k, but as well it requires histograms of regions just exterior to \mathbf{z}_k^j . These may be $N(j,k)$ or sub-regions of these neighbors. Call these exterior regions $N'(j,k)$. Thus we have concluded that $p(\mathbb{I} | \mathbf{z}_{k \cup N(j,k)}^j) = \prod_{m \text{ a subregion of } (j,k)} p(H_m^{j+1}) \prod_{i \in N'(j,k)} p(H_i^{j \text{ or } j+1})$.

7. Results using proposed statistical framework

We have created quite a complex structure. While we have argued each of its steps and might find the structure elegant, what other evidence can we adduce for its validity? As described in the following, object and histogram statistics in our examples have desirable properties. Tries at segmentation by posterior optimization of deformable m-reps using

major parts of this structure are showing success. But far more validation is needed. In particular, we do not have results using all of the proposed framework together, but we have the following results, each using many of its components.

In the companion paper by Pizer, Jeong, Lu, et al. [2005] we have shown that at an object scale level, using a set of three 3D objects $\{\mathbf{z}_k^j\}_k$ in the male pelvis without interstitial regions and using a way to describe $p(\text{shape of } \mathbf{z}_k^j, \text{inter-relationship of } \mathbf{z}_k^j \text{ and } \mathbf{z}_{N(j,k)}^j)$ that is explained in that paper, we find that samples from the generated probability distributions are nearly geometrically proper and that the probability distributions have intuitively reasonable means and principal modes of variation.

We have carried out segmentation experiments using a three-scale structure on m-rep models of the three objects: bladder, prostate, and rectum in CT images from a population over many days from the same patient. In these experiments the three scale levels are the global scale, the object scale, and the medial atom, i.e. through-object-subsection, scale. The image contrasts at the boundary of the prostate are extremely low at large portions of its boundary (see Fig. 3). The log likelihood, i.e., $\log p(\mathbb{I} | \mathbf{z}_{k \cup N(j,k)}^j)$ values used were based on normalized intensity correlations rather than our proposed histogram-based probabilities. As illustrated in Fig. 3, on trials on multiple days of one patient, the segmentation appeared to be clinically acceptable [Chaney 2004].

Studies reported in the companion paper by Broadhurst et al. [2005] produced segmentation results of the bladder, prostate, and rectum on the aforementioned patient's CTs for the global scale level only. The results from a variety of methods for evaluating $\log p(\mathbb{I} | \mathbf{z}_{k \cup N(j,k)}^j)$ were compared. These included a variety of means of determining the structure and size of the regions over which histograms were computed. A method using histogram statistics showed improvements over the method using normalized correlation to measure $\log p(\mathbb{I} | \mathbf{z}_{k \cup N(j,k)}^j)$.




Fig. 3. Slice of CT of male pelvis and 3D segmentation result using a multi-scale posterior optimization of an m-rep, largely according to the schema described in this paper. Shown in white is a cross section of a manual segmentation, for comparison with the computer segmentation shown in color.

8. Discussion

The method of histogram probabilities has already been extended to populations \underline{I} describing a collection of values at each pixel or voxel, where the values may be separate such as luminance and two chromanences, or they may be derived from \underline{I} , such as a collection of derivatives of \underline{I} or a collection texture features, e.g., obtained by Gabor filtering with different filters [Broadhurst 2004].

We have much work left to do. The details of many of the stages of this schema are left to be designed, implemented and evaluated. Among the open questions are how precisely should the regions be formed; what is the best way to represent inter-region geometric relationships; what is the best way to produce statistics on these inter-region geometric relationships; how precisely do we transition from the large and moderate scales at which m-reps seem a particularly attractive object representation to the scale of the individual voxel; how do we transition from histograms at these large and moderate scales to Gaussian scale space at the scale of the individual voxel.

Yet the initial results are quite attractive. They certainly overcome the problem associated with regional summaries into a single intensity, that intensities are located more precisely than their scale would suggest, when generating probabilities on image intensities. This problem is found in not only Gaussian scale spaces but also those produced from geometry-limited diffusion. We think our multiscale schema for working with deep structure of images in populations via geometric models in populations is worth following and recommend it to others for consideration and improvement.

Acknowledgements

We are grateful to Sarang Joshi and Edward Chaney for scientific discussions and to Delphine Bull for help in preparing this paper. The work described in this paper was done under the partial support of NIH grant P01 EB02779.

References

Broadhurst, RE (2004). Simplifying Texture Classification. Univ. of N.C. Dept. of Computer Science technical report TR05-009. http://midag.cs.unc.edu/pubs/tech-rpts/Broadhurst_TR05_9.pdf

Broadhurst, RE, J Stough, SM Pizer, E Chaney (2005). Histogram Statistics of Local Model-Relative Image Regions. In this volume, *International Workshop on Deep Structure, Singularities and Computer Vision*.

Chaney, E, S Pizer, S Joshi, R Broadhurst, T Fletcher, G Gash, Q Han, JY Jeong, C Lu, D Merck, J Stough, G Tracton, MD J Bechtel, J Rosenman, YY Chi, and K Muller (2004). Automatic Male Pelvis Segmentation from CT Images via Statistically Trained Multi-Object Deformable M-rep Models, Presented at American Society for Therapeutic Radiology and Oncology (ASTRO).

- Cootes, TF, GJ Edwards and CJ Taylor (1998). Active Appearance Models. *Proc. European Conference on Computer Vision 1998* (H.Burkhardt & B. Neumann Ed.s). **2**: 484-498, Springer.
- Fletcher, PT, C Lu, SM Pizer, S Joshi (2004). Principal Geodesic Analysis for the Study of Nonlinear Statistics of Shape. *IEEE Transactions on Medical Imaging*, **23**(8): 995-1005.
- Florack, L, A Kuijper (2000). The Topological Structure of Scale-Space Images. *J. Mathematical Imaging and Vision*, **12**(1): 65-79.
- Grenander, U (1976). *Pattern Synthesis: Lectures in Pattern Theory*, volume I. Springer-Verlag.
- Grenander, U (1978). *Pattern Synthesis: Lectures in Pattern Theory*, volume II. Springer-Verlag.
- Han, Q, SM Pizer, D Merck, S Joshi, JY Jeong (2005). Multi-figure Anatomical Objects for Shape Statistics. To appear, *Information Processing in Medical Imaging* (IPMI). Springer LNCS.
- Ho, S. (2004). Profile Scale Spaces for Statistical Image Match in Bayesian Segmentation. Dissertation, Univ. of N.C. Dept. of Computer Science, <http://midag.cs.unc.edu/pubs/phd-thesis/SHo04.pdf>
- Kendall, DG (1989). A survey of the statistical theory of shape. *Statistical Science*, **4**(2): 87-120.
- Lindeberg, T (1994). Scale-space theory: a basic tool for analyzing structures at different scales. *Journal of Applied Statistics: Special issue on Statistics and Images*, **21**(2): 223-261.
- Lu, C, SM Pizer, S Joshi, JY Jeong (2004). Statistical Multi-object Shape Models. In preparation. <http://midag.cs.unc.edu/LuC2004.pdf>
- Pennec, X (1999). Probabilities and statistics on Riemannian manifolds: basic tools for geometric measurements. In *IEEE Workshop on Nonlinear Signal and Image Processing*.
- Pizer, SM, JY Jeong, C Lu, K Muller, S Joshi (2005). Estimating the Statistics of Multi-Object Anatomic Geometry Using Inter-Object Relationships. In this volume, *International Workshop on Deep Structure, Singularities and Computer Vision*.
- ter Haar Romeny, BM (2003). *Front-End Vision and Multi-Scale Image Analysis: Multi-Scale Computer Vision Theory and Applications, written in Mathematica*. Kluwer Academic.