Contents

Chapter 1

Introduction

Recent years have witnessed a tremendous growth in the role of computing in image-guided Medicine and Neuroscience. Interventional radiologists and surgeons at leading research institutions have upgraded the "light-box method" of studying images in which one views films of two-dimensional cross-sections of a patient's medical scan by mounting them on a light-box (see Figure 1-1), to one where they now routinely use three-dimensional visualizations of the data such as the ones shown in Figures 1-2,1-3, (See [ETH, MAYO, MNI, SPL, MGH, JHU, SHADYSIDE, LEUVEN, UCLA, UCSF, GUYS, GE, MERL, INRIA] for overview of image guided medicine at some of the leading research laboratories).

Many facets of image-guided Medicine have benefited from increased use of computing in processing images. In a large number of these applications, the "segmentation" problem needs to be addressed, and is often considered a bottleneck. The focus of this thesis is segmentation of medical imagery. Our contribution is a novel method for segmentation of medical images that combines knowledge of the gray-level appearance of different tissues, the characteristics of the imaging modality that is used to acquire the images, and knowledge of spatial relationships between anatomical structures.

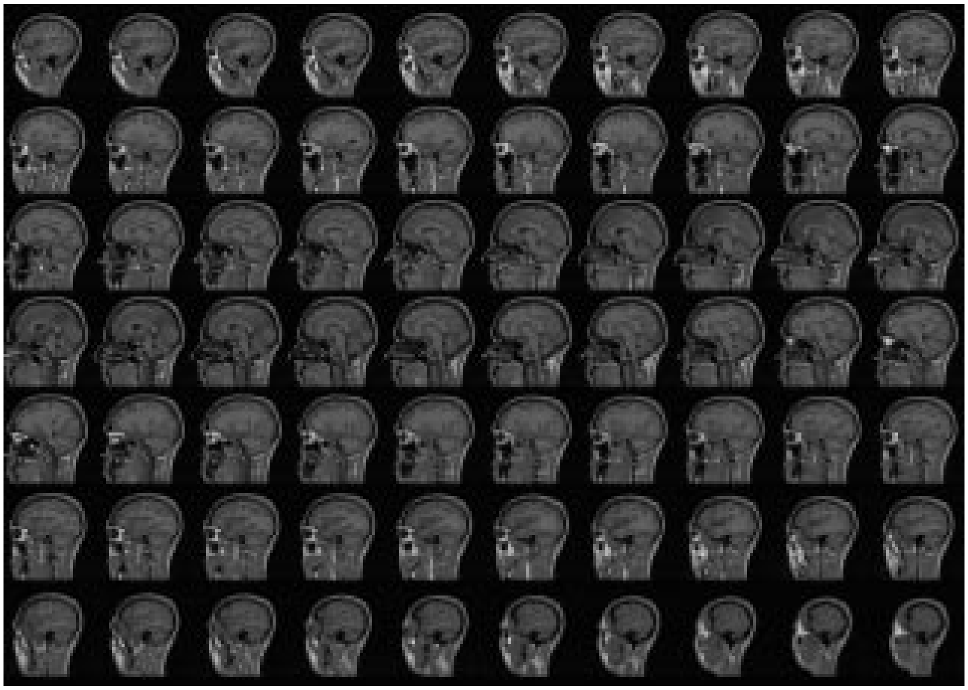

Figure 1-1: A set of two-dimensional sagittal cross-sectional slices comprising a 3D brain MRI scan. This is the view of an MRI scan that is printed on film and mounted on a light-box and is traditionally available to clinicians to aid them in making diagnosis and planning therapy.

1.1 The Medical Image Segmentation Problem

Segmentation of digital medical imagery is a labeling problem in which the goal is to assign to each voxel[1] in an input gray-level image, a unique label which represents an anatomical structure. This labeled image is referred to as the "segmentation" of the input image. An example input gray-scale image is a Magnetic Resonance Imaging (MRI) scan of a human head that was shown in Figure 1-1. A labeling or segmentation for one such cross-sectional image into skin, brain tissue, ventricles, and tumor, and a three-dimensional visualization created from segmentation of a complete scan are shown in Figure 1-4.

In a different example, Figure 1-5 shows three-dimensional visualizations of the bones, the cartilage, and the ligaments of the human knee as well as the segmented images from which they are constructed.

---

[1] A voxel is a volume element of a three-dimensional image. It is the three-dimensional equivalent of a pixel in a two-dimensional image.

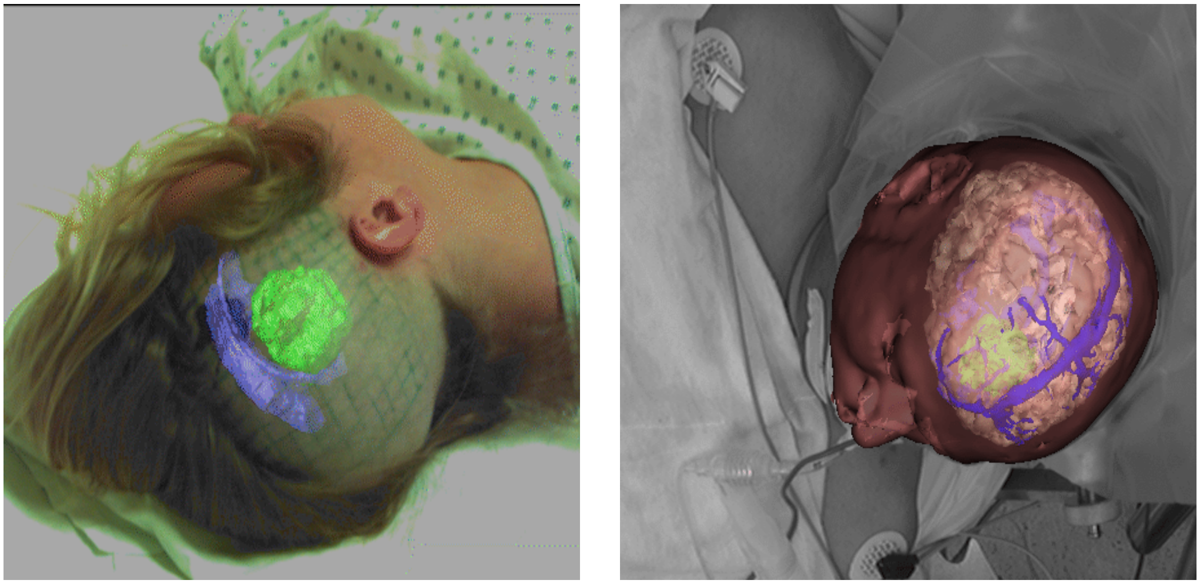

Figure 1-2: Three-dimensional visualization used by Neurosurgeons to examine the relationship of pathology to familiar landmarks. These visualizations are used for pre-operative surgical planning as well as intra-operative guidance and verification.

1.2 Applications of Segmentation in Medicine

Image-Guided Surgery: The most visible, and perhaps highest impact, use of segmented images has been in image-guided surgery [Grimson et al.1997, Ayache1995, Colchester et al.1996, Bucholtz et al.1997] where frequently the goal is to provide intra-operative *enhanced reality visualization*, such as the ones shown in Figure 1-2, to surgeons.

End-to-end image-guided surgery systems often consists of three modules: one for segmentation, one for "registration", and one for "tracking". The segmentation module consists of an off-line mechanism, such as the one proposed in this thesis, to generate segmentations of key anatomical structures which are rendered in three-dimensions and serve as landmarks during surgery. The registration module provides a mechanism for transforming pre-operative imagery into the intra-operative coordinate-frame of the patient. The tracking module follows the movement of the patient during the surgery and updates the enhanced reality visualization to be consistent with the view of the patient that is visible to the surgeon. This allows surgeons to visualize the location of their surgical instruments relative to the

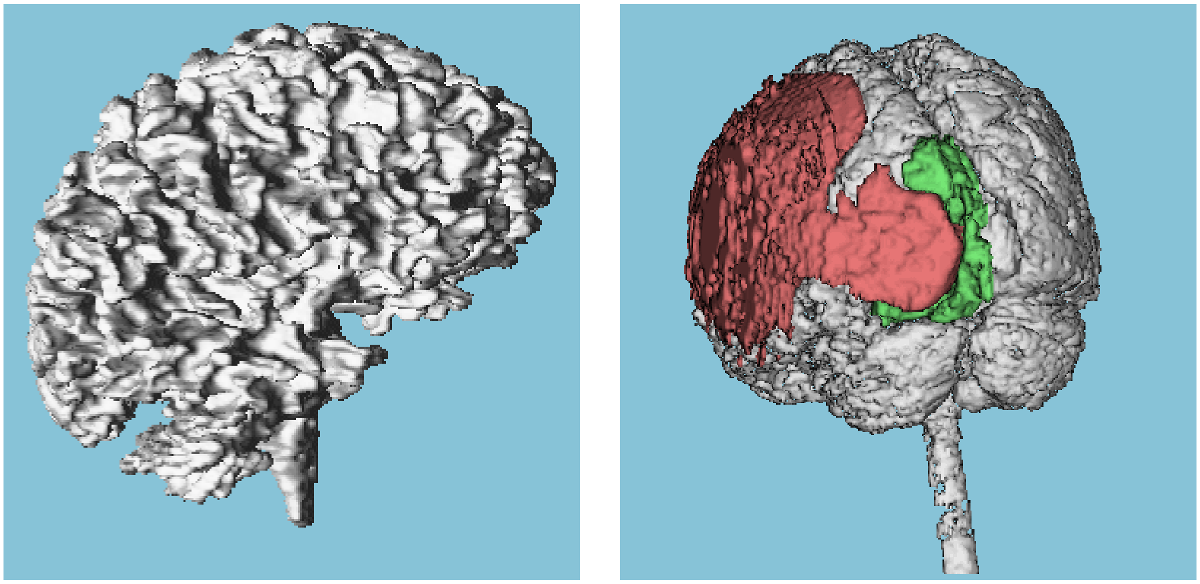

Figure 1-3: Left: a 3D rendering for visualization of the white matter from an MRI scan of a normal subject. Right: a three-dimensional rendering of the brain surface of a patient, overlaid with pathology which is shown in green and red. Such visualizations are now routinely used in leading research institutions for diagnosis, therapy planning and therapy evaluation.

landmarks in the pre-operative scans. State of the art image-guided surgery systems perform a rigid registration with six degrees of freedom – three for rotation, three for translation – to align the pre-operative and intra-operative imagery, and use off-the-shelf tracking hardware that uses either optical or electro-magnetic sensors to track the motion of the patient during surgery. XPlan [Grimson et al.1994, Grimson et al.1995, Black and Chaberie1998], an image-guided surgery system that has been used in over two hundred neurosurgeries at Brigham and Women's hospital in Boston, MA, uses semi-automatic segmentation of pathology and normal tissue, a rigid registration technique [Ettinger1997], and tracks patient motion using an optical tracker [Leventon1997]. The enhanced reality visualizations shown in Figure 1-2 were generated using XPlan.

While improved automation of segmentation is necessary for wider clinical acceptance of image-guided surgery, the most interesting advances in the area are likely to be the ones that will relax the simplifying assumption of rigidity between pre-operative and intra-operative anatomy, and account for the changes in anatomy that take place

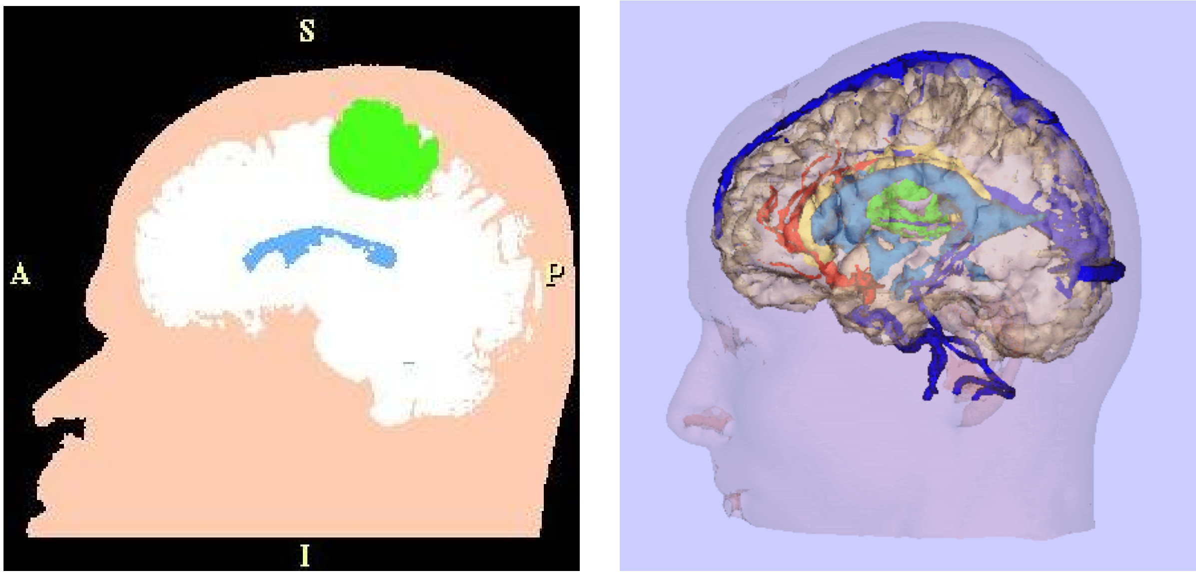

Figure 1-4: Left: a segmented MRI slice showing skin in pink, brain tissue in white, ventricles in blue and tumor in green. Right: a three-dimensional rendering of segmented skin surface, brain tissue, major blood vessels, and tumor from a diagnostic MRI scan.

as surgery progresses. The availability of interventional Magnetic Resonance Imaging [Jolesz et al.1997] and companion tools such as the 3D Slicer [Gering1999] now allow surgeons to capture high contrast intra-operative images of the changes in the positions of anatomical structures such as the swelling of the brain when the skull is opened, the egression of cerebrospinal fluid, or resection of tumor tissue during surgery. In the coming years this ability will almost certainly be instrumental in the development of computational, predictive models of deformation that will allow us to reconcile pre-operative, intra-operative and post-operative anatomy in a way that is not available to us today.

Surgical Simulation: In a different surgical domain, researchers are building surgical simulators to provide surgeons with virtual testbeds for planning and rehearsing patient specific procedures, with the motivation that such preparation will make the actual procedures on patients more efficient (rehearsed surgeons will work faster) and effective (rehearsed surgeons will produce better results for the patient) [O'Toole1998]. As is the case in image-guided surgery systems, segmentation is only one component in surgical simulation; for example, a surgical simulator for arthro-

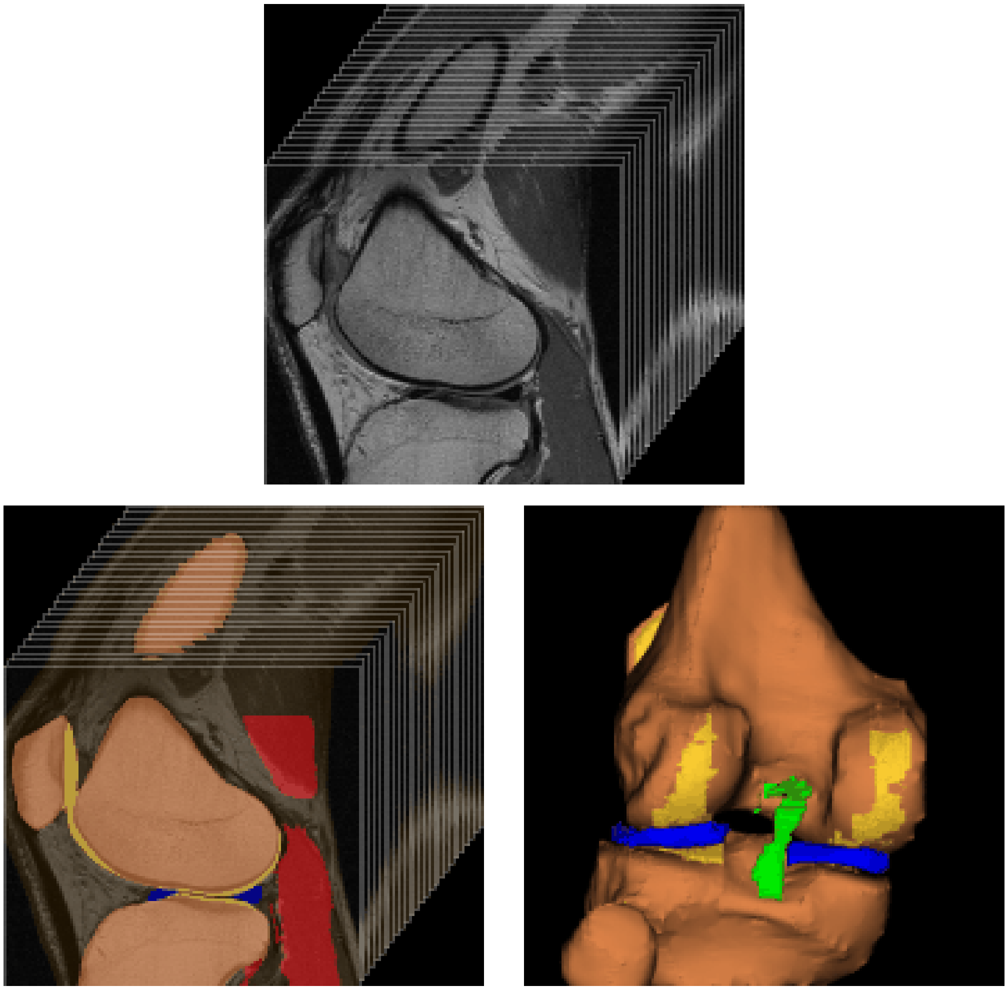

Figure 1-5: Top: A stack of 2D cross-sectional slices comprising a 3D knee MRI scan. Bottom Left: Segmentation as a stack of 2D slices showing bones, cartilage, and ligaments. Bottom Right: A 3D rendering of the individual structures in the segmented image.

scopic knee procedures that is underway at Mitsubishi Electric [Gibson et al.1998] uses the visualizations generated from segmentation of pre-operative knee MRI in conjunction with physics-based models of structures, and force-reflective interface devices to create a complete simulation environment. While segmentation is the focus of this dissertation, other components of this simulator have lead to the development of ground-breaking techniques for measuring the mechanical stiffness of tissue, specifically articular cartilage [Samosky1998], which when incorporated into the computer models of the knee, promise to increase the accuracy and realism of arthroscopic surgical simulation.

Neuroscience Studies and Therapy Evaluation: Dominant non-surgical uses of segmentation include Neuroscience studies and drug therapy evaluation. A large class of Neuroscience studies analyze segmented medical images of several subjects in different populations such as “normals” or “schizophrenics”, with the goal of being able to characterize the anatomical features that distinguish the populations. For example, a recent study in schizophrenia determined the size of the lower left temporal lobes as a differentiating characteristic between normal and schizophrenic subjects [Hirayasu et al.1998]. Other studies seek to determine the role of not just the sizes of structures, but their shape, in characterizing a population [Martin1995].

Another high-impact application domain for segmentation is drug therapy evaluation in which normal and pathological tissue are segmented in registered scans that are taken at different time points in the duration of the therapy, and the changes in the characteristics of pathology are used to measure the effect of the drugs that are under evaluation. Efficient automatic processing of images to extract this drug-evaluation information can shorten the time it takes to bring potentially life-saving drugs to market. Recently, a large therapy evaluation trial for Multiple Sclerosis was conducted where the volume of white matter lesions was measured from a time-series of segmented brain MRI scans to evaluate the impact of a particular drug [Warfield and Kikinis1998].

In all these applications, segmentation of input imagery is a pre-processing step which is often considered the bottleneck in the application.

1.3 The Difficulties in Medical Image Segmentation

Limitations in Manual Segmentation: Given the compelling applications of segmentation, it is natural to first explore a manual solution to the problem, i.e., one in which an expert anatomist manually delineates the boundaries of different structures in medical images. Such a solution is tedious; a typical scan which shows good contrast between soft tissue consists of approximately eight million voxels across a hundred cross-sectional images that need to be classified. Moreover, studies have shown that manual segmentations vary significantly between experts as well as when an expert segments the same image at different times. Up to 15% variability has been reported in the segmentation of cortical gray matter of the brain in segmentations by five different experts [Warfield and Kikinis1998], and between $15 - 22\%$ variability has been reported in segmentation of brain tumors by multiple experts [Kaus et al.1998]. These variations render the manual solution problematic in tasks such as the Neuroscience studies mentioned above where the desired accuracy is better than the variability between manual segmentations, and underscore the need for reproducible and accurate automatic segmentation methods.

Challenges in Automatic Segmentation: In automatic segmentation, reproducibility is less of a concern because most algorithms that are used are deterministic; the challenge lies in effective representation and use of domain knowledge – knowledge that is acquired by human experts through extensive medical training – about the underlying anatomy and the different imaging processes that are used to acquire the scans. This representation and use of domain knowledge is challenging simply because the underlying anatomy that is to be captured is complex in shape and appearance, and varies across subjects. Complexity in shape here refers to the fact that anatomical structures are not regular geometric shapes such as circles, rectangles or straight lines that are easy to represent computationally. Complexity in appearance refers to the fact that anatomically distinct tissues can have the same gray-level appearance in an image due to the characteristics of the imaging modality (for example, in gradient echo MRI, two distinct tissues: the scalp and white matter, are both mapped into overlapping medium gray brightness), while the same tissue can have different gray-level appearance in different regions of the image due to distortions caused by the imaging equipment (white matter in one region of the image can be considerably brighter than white matter in another region due to the inhomogeneity of the radio frequency (RF) coils in MRI).

1.4 Segmentation Approach in this Thesis

The segmentation approach proposed in this thesis is based on the belief that knowledge of anatomy and the imaging process is crucial to the success of segmentation. Specifically, a segmentation method for a given structure must have sufficient information about the domain to be able to answer the following four questions:

- What is the *intensity* or gray-level appearance of the structure in the image?
- What are the characteristics of the *imaging modality* that was used to acquire the image that impact the gray-level appearance of different structures in the image?
- What is the *shape* of the structure?
- What is the *geometric relationship* of the structure to other structures in the image?

We propose a framework the medical image segmentation problem is posed as a knowledge-based task that uses four types of models, one for each piece of domain knowledge mentioned above. These models are:

- *Intensity models* that describe the gray level appearance of individual structures. These encode information such as "fluid appears bright in T2-weighted MRI".
- *Imaging models* that describe characteristics of the medical imaging modality that is used to create the imagery. These encode information such as the presence of inhomogeneity in MRI images due to limitations in the imaging equipment.
- *Shape models* that describe the shape of structure in a given population. These encode information such as "the ventricles are bent in the middle in healthy adults".

- *Geometric models* that describe the spatial relationships between structures in a given population. These encode information such as "articular cartilage lines the femur in healthy adults".

Automatic segmentation algorithms frequently use some form of *intensity models* to describe the gray-level appearance of anatomical structures. Typically used intensity models are in the form of probability distributions on individual voxels that belong to a particular structure. For example, the intensities of white and gray matter in brain MRI have been represented using the Gaussian distribution [Vannier et al.1985, Gerig et al.1989, Cline et al.1990, Chakraborty, Staib and Duncan1994, Wells III et al.1994] and also as mixtures of Gaussian distributions [Aylward et al.1994]. Richer intensity models capture higher order correlations in gray-levels i.e. describe the appearance of local patches of the structure instead of individual voxels. For example, the mean and variance of local patches are used to characterize the appearance of trabecular bone in MRI [Lorigo et al.1998], and tensors are used to describe the appearance of thin bone in CT images [Westin and Kikinis1998].

The use of explicit *imaging models* to account for the distortions in medical imagery that are attributable to limitations of different modalities such as MRI, CT, PET, SPECT, XRay, ultrasound, is not very common in automatic segmentation algorithms. Segmentation methods for MRI, CT and XRay images that incorporate knowledge of the imaging modality report significant improvements in their results compared to methods that do not account for the distortions in the signal that are due to the characteristics of the imaging equipment [Wells III et al.1996, Highnam1992].

Shape models are an actively researched area in medical image understanding. The idea behind such computational models is to capture salient characteristics that define the "shape" of a structure across a population. Note that this definition is frequently circular and relies on our intuitive understanding of the concept of shape. While parametric models have been developed for individual applications [Staib and Duncan1992, Cootes et al.1994], there is not yet a satisfactory emergent definition of shape that can be applied to anatomical structures in general. If available, shape models can be used to bias a segmentation algorithm towards a particular shape, especially when the gray-level appearance isn't sufficient to uniquely segment the structure. A style of shape modeling that has dominated the literature recently [Szekely et al.1996, Solloway et al.1997] derives from a combination of *deformable models* and *principal component analysis*(PCA). Deformable models are variants of the classical "snake" that was originally proposed in [Kass, Witkin and Terzopoulos1988]. They are used in segmentation to incorporate constraints of local smoothness (see [McInerny and Terzopoulos1996] for a recent survey of deformable models). Principle components capture the directions of largest variation in a given set of points, and are used to characterize the space of possible deformations for a structure in a given population (see textbooks [Duda and Hart1973, Bishop1995, Therrien1984] for a discussion of PCA). In the PCA-driven deformable model style of shape modeling, the models are represented using principle components of normalized spatial locations of feature points that lie on a structure, from a set of training images for a given population. Given a novel image, a deformable model is initialized with the mean shape from the training data. The deformable model is evolved under local constraints of smoothness, and global constraints that require that its shape remains within the space restricted by the principle components. The efficacy of PCA driven shape modeling remains to be demonstrated, but it has brought to the foreground two important issues in shape modelings: how should feature points be selected if PCA is to be used to capture the variations in a population, and how should a novel image be normalized with the shape model so that the shape model can be used. These issues should be considered in future research in order to develop a successful strategy for modeling shape.

Geometric Models are a class of models that we introduce in this thesis for the purpose of encoding notions of spatial relationships between structures in an image. The motivating observation for these models is that while some structures can be segmented from medical images by using methods from low-level image processing, e.g., skin surface is reproducibly segmented from head MRI using a combination of thresholding, connectivity, and morphological operations, there are other structures, such as the brain tissue in head MRI, that do not have as salient a combination of intensity and topology as the skin, and are harder to segment using low-level methods. Geometric models facilitate a “coarse to fine” strategy for segmentation in which the easily identifiable or coarse structures are first segmented automatically and their geometry is then used to bootstrap the segmentation of other fine structures in the image.

1.5 Contributions of this Thesis

In this thesis we instantiate the above framework in a novel method for segmentation of MR images that uses *intensity models* to encode the gray-level appearance of different tissues in the image, *imaging models* to account for thermal noise inherent in MR and as well as corruption of the signal due to inhomogeneities in the RF coil, and *geometric models* to encode the spatial relationships between anatomical structures in the image.

1.5.1 Situating the Work Relative to EM-Segmentation

In the proposed method we build upon a Bayesian method for segmentation of MRI, EM-segmentation, that was introduced by Wells et al. in [Wells III et al.1994, Wells III et al.1996]. The key contribution of EM-segmentation was that it explicitly modeled the corrupting field that is due to the inhomogeneities in the RF coil, and used a non-linear optimization technique, Expectation-Maximization, to iteratively estimate the classification (segmentation) of the input image and the corrupting field. EM-segmentation used Gaussian distributions to model the intensities of the different tissues, and a slowly-varying model for the gain field to account for the inhomogeneities in the data. Typical use of the EM-segmentation algorithm involves embedding it in a segmentation pipeline of three steps:

- *Thermal Noise Removal:* Clinical MR scans typically contain thermal noise, and the signal to noise ratio is inversely related to the acquisition time of the scan. While a scan acquired over an hour would be “cleaner” than one acquired in less time, clinical diagnostic scans are usually acquired in less than twenty minutes. This scan-period reflects a tradeoff in scan quality, cost to the insurance company and patient discomfort: clinicians prefer noise-free scans that would take an hour, insurance companies prefer shorter scans to minimize cost, and patients prefer shorter scans because it minimizes the amount of time they spend enclosed in the imaging equipment trying not to move. As a consequence, removal of thermal noise is the first step in the segmentation pipeline, and is often performed using a gradient-limited diffusion filter [Perona, Shiotan and Malik1994, Whitaker and Gerig1994] which is a smoothing filter that preserves edges in the image.

- *EM-segmentation:* The EM-segmentation algorithm is used to correct for the inhomogeneities in the image and produce an intensity based classification of the voxels in the image into different tissue classes.

- *Post-processing:* Geometric information is incorporated by semi-automatic post-processing of the intensity classification. This post-processing is necessary because anatomically distinct tissues can have the same gray-level appearance in an MR image, and geometric information is needed to disambiguate them. For example, in gradient echo MRI the scalp and white matter, which are two anatomically distinct structures, are both mapped into overlapping medium gray brightness in the image, which leads to identical tissue classification for them when only intensities are used. The exact type of post-processing depends on the data. For example, morphological erosion and dilation are frequently used to isolate the brain in intensity classified MR images [Brummer et al.1993].

1.5.2 Additions to EM-segmentation

Our method replaces this pipelined segmentation of MRI with an integrated method that incorporates thermal noise removal and geometric post-processing into the Expectation-Maximization classification mechanism of EM-segmentation. Given the Bayesian nature of EM-segmentation, we formulate our contributions as two *priors* for EM-segmentation:

- *De-Noising Prior:* We introduce a Gibbs prior on tissue class to model the piecewise homogeneous nature of tissue, which results in resistance to thermal noise in the images. We solve this Gibbs model using a Mean-Field (MF) approximation in conjunction with Expectation-Maximization. Gibbs priors have been used in Statistical Physics, Computer Vision and other fields to model piecewise homogeneity of materials [Geman and Geman1984, Besag1986]. Mean-Field approximations to Gibbs models have been used previously as well [Parisi1988, Elfade1993, Geiger and Girosi1991]. The novelty of our de-noising prior is in its use of the Mean-Field Approximation in conjunction with Expectation-Maximization. We refer to this method of using a Mean-Field approximation in conjunction with Expectation-Maximization as EM-MF segmentation.
- *Geometric Prior* We introduce a novel statistical description of the spatial relationships between structures and refer to those as *geometric models.* These are computational models for notions of spatial relationships such as "cartilage lines the bone" or "white matter lies inside the skin boundary", that are likely used by experts in segmentation of structures that exhibit poor intensity contrast in images. We use these models as spatially varying priors on structures in conjunction with EM-segmentation.

We demonstrate the generalizability of our segmentation approach by applying it to two different segmentation problems: segmentation of white matter, gray matter, fluid, and skin/scalp in brain MRI, and segmentation of the bones and the cartilage in knee MRI.

1.6 Roadmap

In the next chapter we present background on Gibbs Random Fields (also known as Markov Random Fields) and provide details of the formulation of the de-noising

Gibbs prior, its Mean-Field Approximation, and how it fits in with Expectation-Maximization. In Chapter 3, we introduce *Geometric Models* and show how they can be constructed from segmented training images, and used as spatially varying priors that encode geometric relationships between structures. In Chapter 4, we present experiments to illustrate the results of using EM-MF segmentation with Geometric Priors on brain MRI and on knee MRI. In Chapter 5 we conclude with suggestions for future work.

Chapter 2

EM-MF Segmentation: Gibbs Modeling in EM-Segmentation

In this chapter we introduce a novel method for segmentation of MR images that we call EM-MF segmentation. This method incorporates Gibbs modeling into EM-segmentation [Wells III et al.1996], which is a well-known method for segmentation of MR images. Gibbs models provide a mechanism for specifying constraints on the interactions between local neighborhoods in an image. Ising used these in Statistical Physics to explain certain empirically observed facts about the the piece-wise homogeneous nature of Ferromagnetic materials [Ising1925]. The classic use of Gibbs models in Computer Vision was for the reconstruction of noisy graylevel images [Geman and Geman1984]. We use them in EM-MF to model the piece-wise homogeneous nature of tissue, and to thus make the segmentation robust to thermal noise in MR.

EM-segmentation is an intensity-based segmentation method that corrects for a slowly-varying gain field that corrupts MR images. By using a Gibbs model in conjunction with EM-segmentation, EM-MF incorporates noise-resistance and intensity-based classification into a single segmentation method. This integrated method can be used as replacement for the first two steps in the three-step pipeline of noise-reduction, intensity-based classification, and geometric post-processing, that is typically used for segmentation of MRI (see Section 1.5.1 for a discussion of this pipeline).

In the rest of this chapter we provide the reader with background material on Gibbs models, the EM-segmentation method, and then describe the EM-MF segmentation method. In Section 2.1 we define the concept of a Markov Random Field and state an important theorem that establishes equivalence between Markov random fields and the algebraically and computationally more tractable Gibbs Random Fields (GRFs). Using a Gibbs model in image reconstruction amounts to solving an optimization problem, and we give an overview of techniques that are typically used to optimize Gibbs random fields.

In Section 2.2 we provide background on the EM-segmentation algorithm. This includes a description of the estimation method, Expectation-Maximization, that is used in EM-segmentation. In this section we also set up notation that we use later in the chapter.

With this background in place, in Section 2.3 we describe the first contribution of this thesis: our EM-MF segmentation method in which we combine Gibbs models with EM-segmentation to produce a noise-resistant method for intensity-based classification of MR images. In Section 2.4 we relate EM-MF to the classical technique of *Relaxation labeling*. In Section 2.5, we show an example of the performance of the EM-MF segmentation method on noisy brain MRI.

2.1 Background on Markov Random Field Modeling

Informally, a Markov random field is a probabilistic model that is defined on a lattice. It is a collection of random variables such that there is one random variable corresponding to each site of a given lattice. Under this probabilistic model, statistical correlations are allowed between a site and its neighbors. We begin by setting up the vocabulary necessary to formally define an MRF. (The content of this section uses material from [Marroquin1984, Li1995].)

A *probabilistic model* on a collection of random variables specifies how probabilities should be computed for events involving subsets – singletons, pairs, triples, or n-tuples – of the random variables. *Statistical independence* is conceptually and computationally the simplest such probabilistic model. It specifies that the probability of a joint event involving k of the variables be computed as a product of their individual probabilities.

A *lattice* is a set of sites or nodes in a graph. An image with m rows and n columns can be represented as an $m \times n$ rectangular lattice where each site corresponds to a pixel in the image. An $m \times n$ lattice is written as a set of indices: $S = \{(i,j) | 1 \leq i \leq m, 1 \leq j \leq n\}$ or using a single index as: $S = \{i | 1 \leq i \leq m \times n\}$. Figure 2-1 shows a 5×5 lattice with the nodes labeled using double as well as single indices.

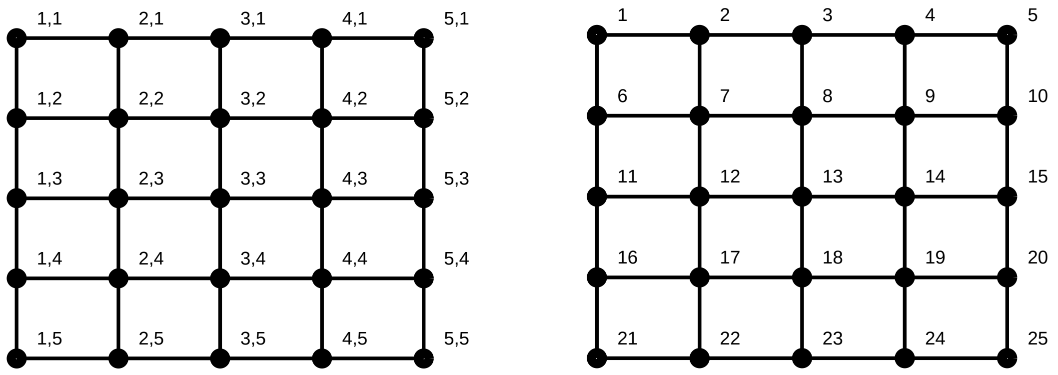

Figure 2-1: A lattice with five rows and five columns. Above and to the right of each node in the lattice is its index, written as a pair in the lattice on the left, and as a single number in the lattice on the right.

A *random field* F is a collection of random variables F_i defined on a lattice S: $F = \{F_i | i \in S\}$.

A *configuration* f of a random field F is a realization or instantiation of F. It assigns values (f_i) to each of the random variables (F_i) in the collection.

The *configuration space* $\mathcal{F}$ of a field F is the set of all its possible configurations.

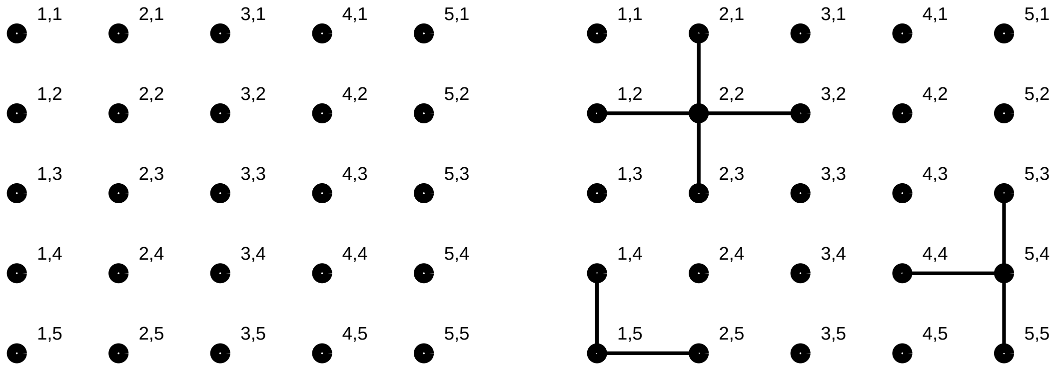

Figure 2-2: A Zeroth order neighborhood system on a lattice is shown on the left. The set of neighbors for each node is the empty set. A first order neighborhood system on is shown on the right. The neighborhood of each internal node (such as node (2,2) in the figure) consists of nodes above, below, to the left and to the right of the node. First order neighbors of edge nodes (such as node (5,4)) and corner nodes (such as node (1,5)) are also shown.

A *neighborhood system* N is a specification of the sites of a lattice S that are the "neighbors" of a site. In a *first order neighborhood system* on an image lattice, as shown in Figure 2-2, the pixels immediately above, below, to the left and to the right of the given pixel are considered its neighbors. The neighbors of an *internal* pixel (i, j) are: $N_{ij} = \{(i-1, j), (i, j-1), (i, j+1), (i+1, j)\}$. Node (2,2) in Figure 2-2 is an example of an internal node with neighbors (2,1), (2,3), (1,2), (3,2). A *second order neighborhood system*, shown in Figure 2-3, includes the first order neighbors as well as pixels that are diagonally across from the given pixel. (A zeroth order neighborhood system, in which the set of neighbors is empty for each pixel, is shown in Figure 2-2.)

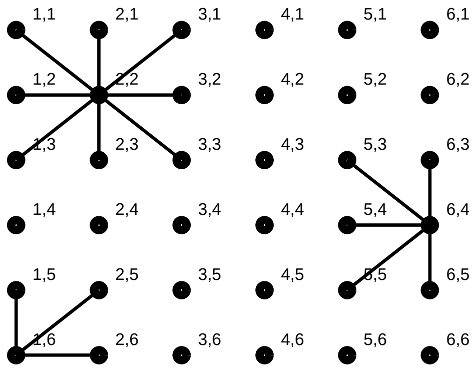

Figure 2-3: A second order neighborhood system is shown here. The neighborhood of each internal node (such as node (2,2) in the figure) consists of its first order neighbors as well as nodes that are diagonally across from it. First and second order neighbors of edge nodes (such as node (6,4)) are also shown.

Definition: Markov Random Field

A random field F is a *Markov random field* on a lattice S with respect to a neighborhood system N if and only if the following two conditions of *positivity* and *Markovianity* are satisfied:

$$P(f) > 0, \forall f \in \{F\} \qquad \text{(positivity)} \qquad (2.1)$$

$$P(f_i | f_{S/\{i\}}) = P(f_i | f_{N_i}) \qquad \text{(Markovianity)} \qquad (2.2)$$

The first equation is more of a technicality in the definition that requires non-zero probability of each configuration of the field. It is the second condition that imposes the Markovianity on the site variables i.e. the condition that the influence of the values of the entire lattice on the distribution of site i is conveyed only by the values of the neighbors of site i. Specifically, it states that the probability of each site variable (f_i) given the state of the rest of the lattice ($f_{S/\{i\}}$, here $S/\{i\}$ is the set of indices in S except the index i, and $f_{S/\{i\}}$ is the set of labels at sites except site i) is computable as the probability of that variable given the state of the site variables in its neighborhood.

In practical applications of Markov random fields, the goal is to find the most probable configuration f^* of a given field F. This optimal configuration f^* is expressed using the argmax notation as follows:

$$f^* = \arg\max_f P(f) \tag{2.3}$$

It turns out that expressing this joint probability has been problematic in Computer Vision applications [Li1995]. In a 1972 result Hammersley and Clifford established equivalence between Markov random fields and Gibbs random fields, which are easier to manipulate algebraically as well as computationally. This MRF-GRF equivalence result is known as the Hammersley-Clifford theorem, and was invaluable for the practical use of MRF theory in Computer Vision (see [Besag1986, Elfade1993, Li1995] for its proof).

2.1.1 Markov Random Fields and Gibbs Random Fields

A random field F is a *Gibbs random field* (GRF) on a lattice S with respect to a neighborhood system N if and only if its configurations satisfy the Gibbs distribution. The Gibbs distribution is commonly used in statistical physics to describe the probability of occurrence of a micro-state f in a system (of molecules) being maintained at temperature T by an external energy source. Intuitively, in a system that obeys the Gibbs distribution, configurations with lower energy are more likely than those with higher energy. Formally, in the Gibbs distribution the probability of a configuration f is given as:

$$P(f) = \frac{1}{Z} \exp\left(-E(f)/\kappa T\right). \tag{2.4}$$

Here T is the temperature of the system. k is the Boltzmann constant. Z is a normalizing constant:

$$Z = \sum_{f'} \exp\left(-E(f')/T\right) \tag{2.5}$$

known as the *partition function.* Division by Z is necessary to make the probabilities sum to one. $E(f)$ is the *energy* of the configuration f, and is expressed as a sum of *clique potentials* $V_c(f)$ over the set C of *cliques* on S:

$$E(f) = \Sigma_{c \in C} V_c(f) \ \ . \tag{2.6}$$

A set of sites in a lattice S is a *clique* c under a neighborhood system N if each site of c is a neighbor of every other site in c under N. The set of cliques of size 1 is denoted C_1, of size 2 is denoted C_2 and so on. The set of all cliques on a lattice (wrt N), C, is the union of all the C_is. Figure 2-4 shows cliques of size 1 and 2 in a first order neighborhood system.

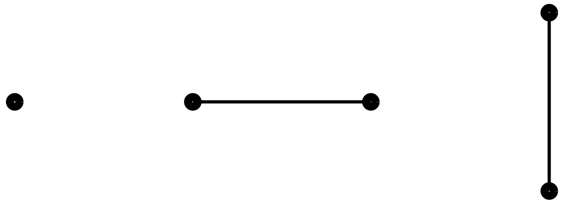

Figure 2-4: The single node on the left indicates the structure of cliques of size 1 (C_1). The other two node-pairs (C_2) show the structure of cliques of size 2 in a first order neighborhood system.

The specification of a GRF includes which cliques should be included in C and what the clique potentials should be.

Theorem: Hammersley-Clifford (1972)

The Hammersley-Clifford theorem states that a field F is a Markov random field on a lattice S with respect to a neighborhood system N if and only if F is a Gibbs random field on S with respect to N.

This theorem allows us to encode interactions between neighboring nodes in a lattice using a Gibbs random field model. Because of the inverse exponential relationship between the energy of a configuration and its probability (see Equation 2.4), the search for the highest probability configuration in a Gibbs random field is a search for the minimum energy configuration f^*:

$$f^* = \arg\min_f E(f) \tag{2.7}$$

Here the optimal configuration f^* occurs at the minimum of the energy function.

2.1.2 Optimization in Gibbs Random Fields

The nature of the energy function $E(f)$ determines which search technique would be most appropriate for finding the optimal configuration f^*. Gradient-descent is a *local* search technique that is efficient to use if the energy function is *convex*. *Global* search methods include stochastic techniques such as *Simulated Annealing* that sample using *Metropolis algorithm* or *Gibbs Sampling*, and deterministic techniques such as *Iterated Conditional Modes*(ICM), *Mean Field Approximation*(MF), and *graduated non-convexity*(GNC). In the following we provide background on these search methods.

Definition: Convex Function

A function $f(x)$ is strictly *convex* over a region R if for all point pairs x_1 and x_2 in R, the following inequality holds:

$$f(\alpha x_1 + (1-\alpha)x_2) < \alpha f(x_1) + (1-\alpha)f(x_2), \quad \forall \alpha \in (0,1) \tag{2.8}$$

The geometric interpretation of convexity is shown in Figure 2-5. A line segment joining any two points on a convex function always lies above the function.

Gradient Descent

Gradient descent is a technique for searching for the minimum of an energy function $E(f)$. Starting from a given initial configuration $f^{(0)}$ of the field, it generates succes-

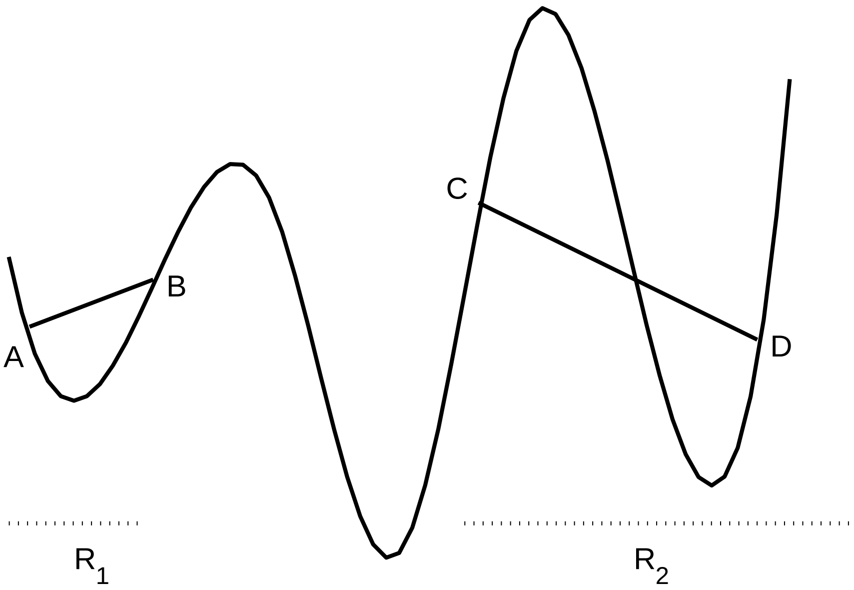

Figure 2-5: A function is strictly *convex* over a given region if the segment joining any two points in that region lies completely above the function. The above function is convex over the region R_1 and non-convex over R_2.

sive configurations $f^{(1)}, f^{(2)}, \ldots$ by moving in the direction of decreasing gradient of the energy function. Given a configuration $f^{(t)}$, configuration $f^{(t+1)}$ is computed as:

$$f^{(t+1)} \leftarrow f^{(t)} - \alpha \nabla E(f^{(t)})$$

where $\nabla E(f^{(t)})$ denotes the gradient of the energy function, and α is a scalar step size that determines the amount of movement in the direction of the decreasing gradient. If $E(f)$ is convex, then under certain conditions gradient descent is guaranteed to converge to its minimum. If $E(f)$ is not convex i.e. has multiple local extrema (as shown in Figure 2-6) gradient descent will converge to the local minimum closest to the starting point of the search.

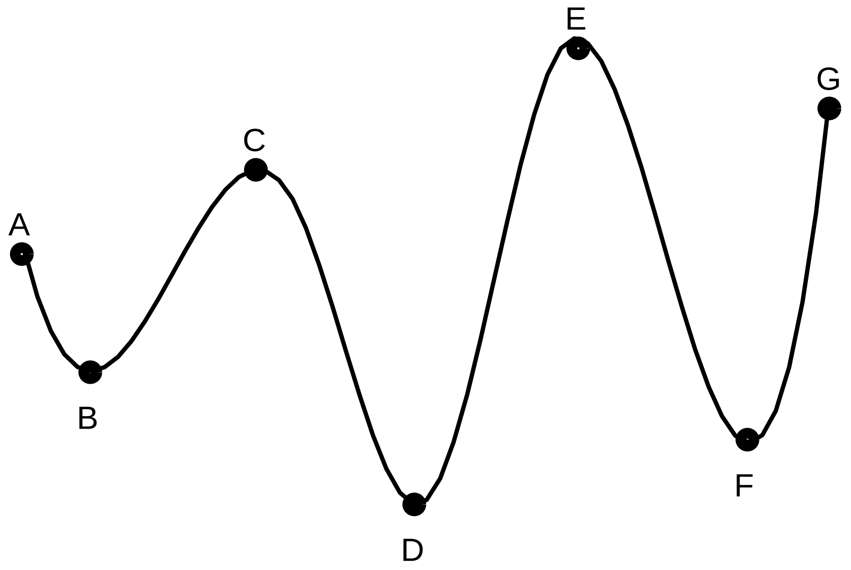

Figure 2-6: B, D, F are local minima, and A, C, G are local maxima of this function.

Simulated Annealing for general $E(f)$

In many problems the energy function $E(f)$ is non-convex. In such cases a search technique called *simulated annealing* can be used to search for the globally minimal energy configuration. Simulated annealing was introduced by Kirkpatrick et al. [Kirpatrick, Gelatt and Vecchi1983].

In metallurgy, slow cooling or *annealing* produces metal that is tougher i.e. has lower energy than what is produced by rapid cooling or *quenching.* The goal of *simulated annealing*, by analogy, is to use a *temperature* parameter to find the global minimum of the energy function, rather than using gradient descent to arrive at a local minimum. The main idea in simulated annealing is to start with a high value for the temperature parameter T, search for a minimum of the energy function raised to the inverse of the T, that is $E(f)^{1/T}$, then lower T and repeat the search until convergence. The search at a fixed temperature is carried out using sampling methods such as the Metropolis algorithm, and the Gibbs sampler which are described later in this chapter. The effect of lowering the temperature, as shown in Figure 2-7, is to make the valleys in the resulting energy function sharper so that in the limit ($T \to 0$) sampling from $E(f)^{1/T}$ will result in configurations that are at the global minimum of $E(f)$.

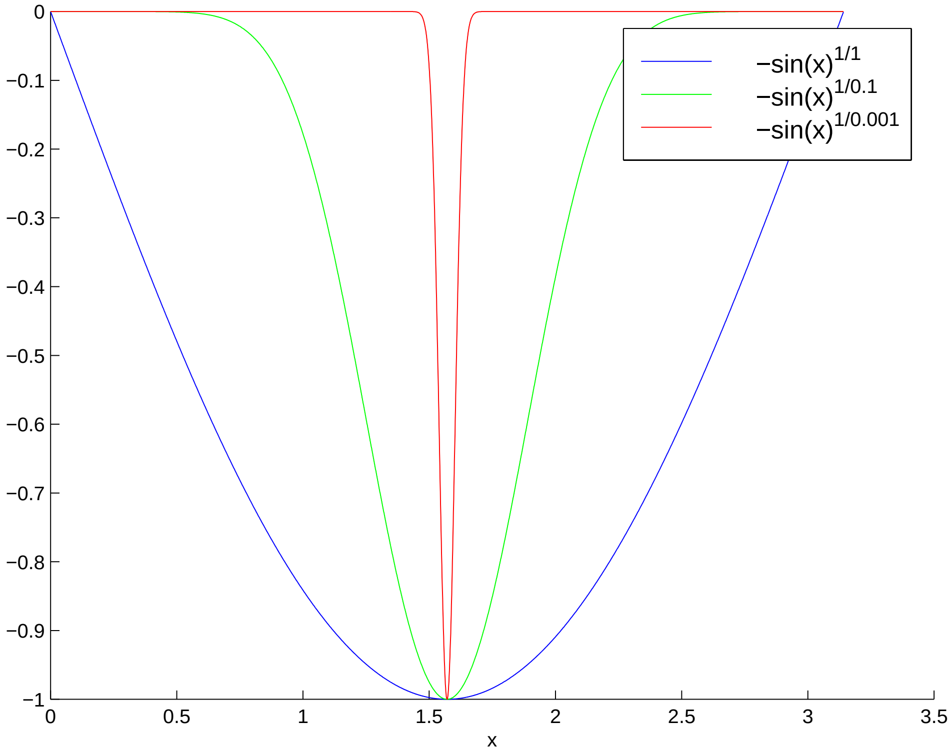

Figure 2-7: Effect of exponentiating a function to the inverse of the temperature T: as T is lowered, the value of the exponent increases, and the extrema in the function are sharpened. Note that the minimum in the above graph gets sharper and approaches a delta function as T is lowered from 1 (blue) to 0.0001 (red).

Sampling Using the Metropolis Algorithm

The Metropolis algorithm [Metropolis et al.1953] is summarized in Table 2.1. It generates a sequence of configurations using a Monte Carlo procedure, and accepts them probabilistically. Starting from an initial configuration, it computes new configurations by small random perturbations of the current configuration. It accepts the new configuration with probability proportional to the change in energy (ΔE) between the old and the proposed states. This probabilistic acceptance criterion allows the algorithm to explore higher energy states (unlike gradient descent) and possibly escape

```
                     $f \leftarrow f^0$                            initialize configuration
repeat:
                     $f' \leftarrow perturb(f)$                    perturb $f$ to get new configuration
                     $\Delta E \leftarrow E(f') - E(f)$            change in energy
                     $p \leftarrow min(1, exp(-\Delta E/T))$       probability of change
                     $r \leftarrow$ random(0,1)                    generate random number
                     if $r < p$, then $f \leftarrow f'$            keep $f'$ with probability $p$
until (convergence)                                                repeat until convergence for fixed $T$
```

Table 2.1: Sampling using the Metropolis algorithm.

local minima.

Gibbs Sampling

Gibbs sampling was introduced in [Geman and Geman1984]. The algorithm is shown in Table 2.2. It computes the next configuration by sampling from the conditional distribution $p(f_i|f_{N_i})$. Gibbs sampling does not use energy functions and is appropriate when the conditional distributions are available.

```
                     $f \leftarrow f^0$                            initialize configuration
repeat:
                     $f' \leftarrow sample(p(f_i|N_i))$            sample from conditional
until (convergence)                                                repeat until convergence
```

Table 2.2: Gibbs Sampling algorithm.

Iterated Conditional Modes

In contrast to the global, stochastic methods described above, Iterated Conditional Modes [Besag1986] is a deterministic, greedy search technique for MRFs/GRFs that converges to a local minimum of the energy. In each iteration, it replaces the discrete label at each site of the lattice with the maximum probability label given the state of its neighborhood. ICM can be interpreted as a special "instantaneous freezing" case of the Metropolis algorithm in which transitions are only allowed to states of lower energy [Li1995].

Graduated Non-Convexity

GNC was introduced in [Blake and Zisserman1987]. It is a deterministic annealing technique in which the energy function $E(f)$ is approximated by a convex function $E^{*}(f)$ which has its single minimum near the global minimum of $E(f)$. If the minima are not close enough, a family of functions “that gradually lose their convexity” are generated from $E^{*}(f)$ to $E(f)$, and the minimum of one function is used to start the search for the minimum of the next one.

Mean-Field Approximation

Mean-field theory provides a framework for developing deterministic algorithms for obtaining approximations to the mean values of random fields. It was originally used in statistical physics [Parisi1988] to approximate the behavior of interacting spin systems in thermal equilibrium. Subsquenty, Mean-Field approximations have been used in Computer Vision and other applications [Geiger and Girosi1991, Peterson and Soderberg1989]. The key step in such approximations involved approximating non-local random variables with their statistical averages. This simplifies the algebraic and computational aspects of the estimation problem. A rigorous development of different MF approximations is presented in [Elfade1993].

This concludes the background material on MRFs/GRFs that is relevant to the understanding of the EM-MF segmentation method. In the next section we provide background material on the EM-Segmentation algorithm, which is useful for understanding EM-MF segmentation.

2.2 Background on EM-Segmentation

[Wells III et al.1994, Wells III et al.1996] introduced a Bayesian method for segmentation of MRI images that is often referred to as EM-segmentation. The key contribution of EM-segmentation is that it explicitly modeled the corrupting field that is due to the inhomogeneities in the RF coil, and used Expectation-Maximization, to iteratively estimate the classification of the input image and the corrupting field.

Other approaches to inhomogeneity correction in MR include [Aylward et al.1994, Pham and Prince1999].

We provide background on Expectation-Maximization in the next section, and review the EM-segmentation method in the following section.

2.2.1 The Expectation-Maximization Algorithm

The Expectation-Maximization algorithm is an estimation method that is used to find the parameters of a model that maximize the likelihood of the data when some some of the data are observable and some are *hidden* or unobserved.

The earliest thorough treatment of the applications of Expectation-Maximization appeared in [Dempster, Laird and Rubin1977]. A recent paper [Neal and Hinton1998] provides an insightful discussion of the algorithm, upon which we base the description of the algorithm in this section.

Starting from an initial guess, Expectation-Maximization re-estimates the model parameters in an iterative manner. Each iteration consists of two steps: an Expectation or E step and a Maximization or M step. In the E-step, the probability distribution for the *hidden* variables is computed, given the known values for the observed variables and the current estimate of the model parameters.

In the M-step, the parameters that maximize the likelihood of the data are re-estimated using assuming that the probability distribution computed for the hidden variables in the E-step is correct. It can be shown that the likelihood does not get worse with each iteration of the EM algorithm.

If we denote the set of observed variables by V, the set of *hidden* variables by H, and the model parameters by θ, the EM algorithm can be described as follows: it starts with an initial guess $\theta^{(0)}$ for the model parameters, and then proceeds to iteratively generate successive estimates $\theta^{(1)}, \theta^{(2)}, \ldots$ by repeatedly applying the following two steps, for $t = 1, 2, \ldots$:

E-Step: Compute $\tilde{P}^{(t)}(H) = P(H|V, \theta^{(t-1)})$. This is the probability distribution for the hidden variables H. It is computed using the current estimate $\theta^{(t-1)}$ of the model parameters and the value V of the observed variables:

M-Step: Compute $\theta^{(t)} = \arg\max_{\theta'} E_{\tilde{P}^{(t)}}[\log P(H, V|\theta')]$. This is expected value of the maximum likelihood (ML) estimate of the model parameters assuming $\tilde{P}^{(t)}$ is the correct distribution for the hidden variables.

The intuition behind the EM algorithm is that if we had the complete data (H, V) – the observable and the hidden variables – we would simply maximize the logarithm of the likelihood $\log P(H, V|\theta)$. But we can not directly compute this log-likelihood because some of our data is missing, and instead we approximate $\log P(H, V|\theta)$ with its *expected value* given the values of the variables that are observable, and a supplied estimate of the model parameters θ_t. It can be shown that under reasonable assumptions this iteration converges to a local maximum of the likelihood [Dempster, Laird and Rubin1977].

Next we review the EM-segmentation method, which uses Expectation-Maximization for segmentation of MR images.

2.2.2 EM-Segmentation

EM-segmentation [Wells III et al.1994, Wells III et al.1996] models the observed MRI signal as a product of the true signal generated by the underlying anatomy, and a non-linear gain artifact due to inhomogeneities in the imaging equipment. It log-transform the observed MRI signal, thus converting the multiplicative gain field to an additive bias field. It uses Expectation-Maximization to simultaneously estimate the classification of the image and the corrupting bias field.

We develop the following notation to describe EM-segmentation, and re-use it in the the next section to describe EM-MF segmentation. This notation is slightly different from the one used in [Wells III et al.1994, Wells III et al.1996].

Notation

- $\mathcal{I}$ is the input image

- N is the number of voxels in $\mathcal{I}$.
- $\mathcal{S} = \{s | 1 \leq s \leq N\}$ is a set of integer indices into $\mathcal{I}$.
- $Y = \{Y_s | s \in \mathcal{S}\}$ is the set of log transformed intensities of $\mathcal{I}$. Y_s is the logarithm of the observed intensity at the s^{th} voxel of $\mathcal{I}$.
- M is the number of classes in the segmentation problem.
- W_{sj} is a binary indicator variable that indicates the membership of a voxel s to class j. It is defined as

  $$W_{sj} = \begin{cases} 1 & \text{if voxel } s \text{ belongs to class } j \\ 0 & \text{otherwise} \end{cases}$$

  A voxel is allowed to be a member of exactly one class, and thus W_{sj} can be 1 for exactly one value of j for a given voxel s. Formally this constraint is expressed as $\sum_{j=1}^{M} W_{sj} = 1, \forall s \in S$.
- W_s is an M-dimensional indicator column vector whose j^{th} component is the indicator variable W_{sj}: $W_s = [W_{s1}, W_{s2}, \ldots, W_{sM}]^T$. The T in the superscript indicates *transpose* of the vector. Note that the state space of the random vector W_s is the canonical basis set $\mathcal{E}_M = \{e_j | e_j = [\delta_{1j}, \delta_{2j}, \ldots, \delta_{Mj}]^T, 1 \leq j \leq M\}$ for $\mathcal{R}^M$. Here δ_{ij} is the Kronecker delta which is 1 when $i = j$ and 0 otherwise.
- W is an $N \times M$ dimensional matrix of all indicator vectors for voxels in the image $\mathcal{I}$: $W = [W_1 W_2 \ldots, W_N]^T$.
- $\mathcal{W}$ is the state space for W: $\mathcal{W} = \underbrace{\mathcal{E}_M \times \mathcal{E}_M \ldots \times \mathcal{E}_M}_{\text{N times}}$
- $\beta = [\beta_1 \beta_2 \ldots \beta_N]$ is the additive bias field that corrupts the image.

Using the above notation, EM-segmentation is posed as an Expectation-Maximization problem in which

- the observed variables are Y - the log transformed intensities in the image,

- the hidden variables are W - the indicator random variables which are used to represent the tissue classification, and
- the model parameters are β - the (additive) bias field corrupting the image.

EM-segmentation uses the observations Y and simultaneously estimates the optimal values for the bias field β and the corresponding values for the classification W.

Specifically, it starts with a guess $\beta^{(0)}$ for the bias field, and then generate successive estimates $\beta^{(1)}, \beta^{(2)}, \ldots$ by applying the following E and M steps, for $t = 1, 2, \ldots$:

E-Step: Computes $\tilde{P}^{(t)}(W) = P(W|Y, \beta^{(t-1)})$. This is the probability distribution for the tissue classification W and is computed using the current estimate $\beta^{(t-1)}$ of the bias field and the value Y of the observed intensities.

This probability distribution is computed as follows:

- Bayes' rule is used to rewrite the posterior probability on tissue classification $P(W|Y, \beta^{(t)})$ as

$$P(W|Y, \beta^{(t)}) = \frac{P(Y, \beta^{(t)}|W)\ P(W)}{\sum_{W' \in \mathcal{W}} P(Y, \beta^{(t)}|W')\ P(W')} \tag{2.9}$$

Then the definition of conditional probability is used to rewrite $P(Y, \beta^{(t)}|W)$ as $P(Y|\beta^{(t)}, W)\ P(\beta^{(t)}|W)$ and the above equation as:

$$P(W|Y, \beta^{(t)}) = \frac{P(Y|\beta^{(t)}, W)\ P(\beta^{(t)}|W)\ P(W)}{\sum_{W' \in \mathcal{W}} P(Y|\beta^{(t)}, W')\ P(\beta^{(t)}|W')\ P(W')} \tag{2.10}$$

Independence is assumed between the bias field and the hidden variables W and $P(\beta^{(t)}|W)$ is expressed as $P(\beta^{(t)})$. The above above expression is rewritten as:

$$\begin{aligned} P(W|Y, \beta^{(t)}) &= \frac{P(Y|\beta^{(t)}, W)\ P(\beta^{(t)})\ P(W)}{\sum_{W' \in \mathcal{W}} P(Y|\beta^{(t)}, W')\ P(\beta^{(t)})\ P(W')} && (2.11) \\ &= \frac{P(Y|\beta^{(t)}, W)\ P(W)}{\sum_{W' \in \mathcal{W}} P(Y|\beta^{(t)}, W')\ P(W')} && (2.12) \end{aligned}$$

This expression for the probability distribution on W is in terms of a measurement model term $P(Y|\beta^{(t)}, W)$ and a prior probability distribution $P(W)$ on W. The denominator is a normalizing constant. Next we provide details on these measurement model and prior terms used in EM-segmentation.

- The measurement model for the log-transformed intensities Y is assumed to be a collection of independent, identically distributed (i.i.d) random variables:

$$P(Y|\beta^{(t)}, W) = \prod_{s \in S} P(Y_s|\beta_s^{(t)}, W_s). \tag{2.13}$$

  Here $P(Y_s|\beta_s^{(t)}, W_s)$ is the predictive probability of the log-transformed intensity value Y_s given the value of the tissue class W_s and the bias $\beta_s^{(t)}$. It is modeled as a Gaussian distribution:

$$P(Y_s|\beta_s^{(t)}, W_s = e_i) = \frac{1}{\sqrt{2\pi\sigma_i^2}}\, exp\ (- \frac{(Y_s - \beta_s^{(t)} - \mu_i)^2}{2\sigma_i^2}) \tag{2.14}$$

  This states that the predictive probability of a particular bias compensated intensity $Y_s - \beta_s^{(t)}$ given that the tissue class is i, is distributed according to a Gaussian whose mean is μ_i and variance is σ_i^2.

- The classification W is modeled as a set of statistically independent random variables. This means that the prior $P(W)$ on W is written as the product:

$$P(W) = \prod_{s \in \mathcal{S}} P^{stat}(W_s) \tag{2.15}$$

- The parameters μ_i and σ_i of the measurement models for the different classes, and the priors, are empirically determined and input to the EM-segmentation algorithm.

M-Step: Computes $\beta^{(t)} = \arg\max_{\beta'} E_{\tilde{P}^{(t)}}[\ \log\ P(\beta'|W, Y)]$. If the complete data

(Y, W) were available, it would simply maximize the logarithm of the posterior probability of the model parameters. Since it doesn't have the complete data, it instead maximizes the expectation of the log-probability of the model over the current estimate of the distribution of the hidden variables. This maximum *a posteriori* probability (MAP) estimate of the model parameters assumes that $\tilde{P}^{(t)}$ is the correct distribution for the hidden variables. Note that this is different from the ML optimality criterion mentioned in the definition of Expectation-Maximization algorithm. It can be shown that the Expectation-Maximization is applicable to MAP estimation in addition to ML estimation.

This MAP estimate of the bias field $\beta^{(t)}$ at iteration t is computed as follows:

$$\beta^{(t)} = \arg\max_{\beta'} E_{\tilde{P}^{(t)}}[\log\ P(\beta'|W, Y)] \tag{2.16}$$

The right hand side can be rewritten as follows using Bayes' rule, and the same independence assumptions that were used in the E-Step:

$$\begin{aligned} &= \arg\max_{\beta'} E_{\tilde{P}^{(t)}}[\log\ (\frac{P(Y|\beta', W)}{P(Y)} P(\beta))] \qquad (2.17) \\ &= \arg\max_{\beta'} E_{\tilde{P}^{(t)}}[log\ P(Y|\beta', W) - \log P(Y) + \log\ P(\beta')] \end{aligned}$$

.

The term not involving β' can be dropped from the right hand side because we are maximizing over β'. The independence of the observations can be used to express the logarithm of the product of the probabilities in $P(Y|\beta', W)$ as the sum of the logarithms of the individual site probabilities $P(Y_s|\beta'_s, W_s)$:

$$\begin{aligned} &= \arg\max_{\beta'} E_{\tilde{P}^{(t)}}[\log\ P(Y|\beta', W) + \log\ P(\beta')] \qquad (2.18) \\ &= \arg\max_{\beta'} E_{\tilde{P}^{(t)}}[\log\ P(\beta') + \sum_s \log\ P(Y_s|\beta'_s, W_s)] \end{aligned}$$

.

The expectation is over W, so the β-term can be pulled out of the expectation:

$$= \arg\max_{\beta'} \left[\log\ P(\beta') + \sum_s E_{\tilde{P}^{(t)}}[\log\ P(Y_s|\beta'_s, W_s)] \right] \ . \tag{2.19}$$

After some manipulation, this can be rewritten as

$$= \arg\max_{\beta'} \left[\log\ P(\beta') + \sum_s E_{\tilde{P}^{(t)}}[W_s^T A_s] \right] \ . \tag{2.20}$$

Using the notation:

$$[A_s]_i = \frac{-(Y_s - \beta_s^{(t)} - \mu_i)^2}{2\sigma_i^2} \ . \tag{2.21}$$

A can be pulled out of the expectation since it does not depend on W:

$$\beta^{(t)} = \arg\max_{\beta'} \left[\log\ P(\beta') + \sum_s E_{\tilde{P}^{(t)}}[W_s^T] A_s \right] \ . \tag{2.22}$$

The above can be rewritten using the notation $\overline{W}_s = E_{\tilde{P}^{(t)}}[W_s^T]$.

$$\beta^{(t)} = \arg\max_{\beta'} \left[\log\ P(\beta') + \sum_s \overline{W}_s^T A_s \right] \ . \tag{2.23}$$

The first term on the right hand side is used to model the smoothness of the bias field using a multidimensional Gaussian distribution. Its logarithm is thus a quadratic form in β'. The second term is also quadratic in β'. Hence the entire expression to be maximized in quadratic in β'. Setting the derivative to zero in the maximization leads to a linear system in β' that can be solved for the updated estimate of the bias field as is detailed in [Wells III et al.1996].

2.2.3 Literature in MRI Segmentation

[Held et al.1998] incorporated a Markov Random Field prior into EM-segmentation that was solved using Iterated Conditional Modes. [Vandermeulen et al.1996] used

Markov random fields for segmentation of MRI. Other treatments of segmentation of MR imagery include [Vannier et al.1985, Cline et al.1990, Zeng et al.1998, Xu et al.1999, Bezdek, Hall and Clarke1993, Gohagan et al.1987, Kohn et al.1991].

2.3 EM-MF Segmentation

In this section we describe the EM-MF segmentation method. It replaces the statistically independent prior on tissue class in EM-segmentation (see Equation 2.15) with a Gibbs prior. Using this Gibbs prior in conjunction with EM-segmentation results in a formalism that uses the Mean-Field approximation to obtain a simple update rule for calculating the expectation of the classification W. This expectation, $\overline{W}$, is used in the M-Step that was described for the EM-segmentation algorithm.

Next we examine how the expectation $\overline{W}$ of the hidden variable W can be computed using the observations and a previous estimate $\beta^{(t)}$ of the model parameters. We first outline the key observations that allow us to arrive at this target expectation, and then fill in the details.

Outline of Assumptions and Observations:

- We model the tissue classification W as a Gibbs random field. In other words, we assume that the prior probability on W, $P(W)$ is a Gibbs distribution.
- We observe that because W is a Gibbs random field, and the measurement models are Gaussian, it turns out (as will be shown later) that $W|Y, \beta^{(t)}$, the random variables whose expectation we seek to estimate, also form a Gibbs random field.
- We observe that
  - because the state space of W_s (and $W_s|Y_s, \beta_s^{(t)}$) consists of canonical basis vectors, the i^{th} component of its *expected value* is the same as the probability of the canonical basis vector e_i, and

- because $W|Y,\beta_{(t)}$ is a Gibbs random field, its expected value, $\overline{W}$, can be approximated using Mean-Field theory from Statistical Physics.

Using these observations, the EM-MF segmentation algorithm can be stated more precisely as: EM-MF uses the Expectation-Maximization algorithm to simultaneously estimate the tissue classification and the bias field in MRI images; it models the hidden variables as a Gibbs random field and uses Mean-Field theory to estimate the expectation of their posterior probability distribution.

Next we fill in the details on how the expectation $\overline{W}$ of the posterior probability distribution of the hidden variables W is estimated using the Gibbs prior, and its Mean-Field solution.

The Gibbs Prior:

In contrast to the EM-segmentation algorithm that was described in [Wells III et al.1996] in which the image is modeled as a collection of statistically independent random variables, in EM-MF segmentation we use a model for the image lattice that relaxes the statistical independence assumption between the voxels of the image and allows local spatial interactions between them. Formally we achieve this by modeling the prior on the tissue classification as a Gibbs random field by specifying a first order neighborhood system on the image lattice and an energy function on the configurations of the field that is similar to the well-known Ising Potential. Such models have been used previously in Computer Vision and Physics to model piece-wise homogeneous phenomenon [Geman and Geman1984], [Besag1986], [Marroquin1984], [Li1995]. In the rest of this section we provide details of our Gibbs prior.

We model our *a priori* assumptions about the hidden variables as a Gibbs random field. This allows us to write the prior probability $P(W)$ associated with W as a Gibbs distribution:

$$P(W) = \frac{1}{Z}\,\exp\big(-\frac{1}{T}E_p(W)\big) \tag{2.24}$$

$$Z = \sum_{W'\in\mathcal{W}} \exp\big(-\frac{1}{T}E_p(W')\big), \tag{2.25}$$

where T is a *temperature* parameter[1], Z is a normalizing constant known as the *partition function*, and $E_p(W)$ is the *energy* associated with the field. In order to complete our specification of the Gibbs random field prior, we define a *neighborhood system* $\mathcal{N}$ on our image lattice, and use that to define the energy function itself.

A neighborhood system $\mathcal{N} = \{N_s | N_s \subset S\}$ specifies the scope/locality of the interactions between the neighbors of the Gibbs random field. We limit the neighbors N_s of a site (voxel) s in the (three-dimensional) image to be its six closest neighbors in $\mathcal{R}^3$. If the index s for a voxel corresponds to the 3d location (x, y, z) in the image, we define the neighbors of site s as:

$$N_s = N_{xyz} = \{(x{+}1, y, z), (x{-}1, y, z), (x, y{+}1, z), (x, y{-}1, z), (x, y, z{+}1), (x, y, z{-}1)\} \tag{2.26}$$

Using this neighborhood system, we define the energy of a configuration or realization w of the field as:

$$E_p(W) = -\frac{1}{2} \sum_{s \in \mathcal{S}} \sum_{r \in N_s} W_s^T J W_r - \sum_s f^T W_s \quad . \tag{2.27}$$

Here J is a $M \times M$ dimensional *tissue-class interaction matrix* and is designed to encode the compatibility between different tissue classes as neighbors. We measure J from training data such that the entry J_{kl} is the logarithm of a ratio of frequencies:

$$J_{kl} = \log \ \frac{\#\text{ times class } k \text{ occurs next to class } l \text{ in training data}}{\#\text{ times class } k \text{ occurs in training data}} \tag{2.28}$$

f is a M-dimensional column vector designed to encode the stand-alone probability of each tissue class. We measure it from training data such that the i^{th} component of f, f_i, is the logarithm of the stationary and prior probability p_i^{stat} that a given voxel belongs to tissue class i: $f_i = \log\ p_i^{stat}$, where p_i^{stat} is estimated as the relative frequency of tissue label i in training data:

---

[1] which will be assumed to be 1 if it is not explicitly stated henceforth

$$f_i = \log \frac{\#\text{ times class } i \text{ occurs in training data}}{\#\text{ voxels in training data}} \quad . \tag{2.29}$$

The Gibbs Posterior:

Now we show that because our prior is a Gibbs form with energy $E_p(W)$, and the measurement model is Gaussian, it turns out that the posterior $W|Y, \beta^{(t)}$ is a Gibbs form with energy $E(W)$. In this section we develop the expression for $E(W)$ in terms of the measurement model and the prior energy $E_p(W)$. We use Equations 2.11 and 2.24 and obtain the following:

$$\begin{aligned} P(W|Y, \beta^{(t)}) &= \frac{P(Y|\beta^{(t)}, W) \ \frac{1}{Z} \ \exp\left(-E_p(W)\right)}{\sum_{W' \in \mathcal{W}} P(Y|\beta^{(t)}, W') \ \frac{1}{Z} \ \exp(-E_p(W'))} && (2.30) \\ &= \frac{\exp\ \left(-E_p(W) + \ \log\ P(Y|\beta^{(t)}, W)\right)}{\sum_{W' \in \mathcal{W}} \ \exp\ \left(-E_p(W') + \ \log\ P(Y|\beta^{(t)}, W')\right)} && (2.31) \end{aligned}$$

This can be expressed as a Gibbs form:

$$P(W|Y, \beta^{(t)}) = \frac{1}{Z'} \ \exp\left(-E(W)\right), \ Z' = \sum_{W' \in \mathcal{W}} \ \exp\left(-E(W')\right) \tag{2.32}$$

where the energy $E(W)$ is given by:

$$\begin{aligned} E(W) &= E_p(W) - \ \log\ P(Y|\beta^{(t)}, W) && (2.33) \\ &= -\frac{1}{2} \sum_{s \in \mathcal{S}} \sum_{r \in N_s} W_s^T J W_r - \sum_s g_s^T W_s \quad . && (2.34) \end{aligned}$$

The above expression is obtained by substituting the definition of $E_p(W)$ from Equation 2.27 and by defining g is as an $N \times M$ dimensional matrix whose si^{th} entry is given by the logarithm of the product of the stationary prior probability of tissue class i and the measurement model at voxel s for tissue class i:

$$g_{si} = f_i + \ \log\ P(Y_s|W_s = e_i, \beta_s^{(t)}) \quad . \tag{2.35}$$

Next we present the Mean-Field solution to the GRF for $W|Y, \beta^{(t)}$, and note that that can be used as an estimate of the target distribution.

The Mean-Field Solution:

Mean-Field theory from Statistical Physics allows us to approximate the *expected value* of a Gibbs random field as the solution to a set of consistency equations. This solution may be obtained as a fixed point of an iteration. For the specific Gibbs random field defined in Equation 2.32, the expected value of $W|Y, \beta^{(t)}$, or its mean-field solution, can be computed as the value of $\overline{W}$ that results from the following iteration (one for each $\overline{W}_{si}$) to convergence:

$$\overline{W}_{si} \leftarrow \frac{\exp\ \left(g_{si} + \sum_{s' \in N_s} \sum_{k=1}^{M} J_{ik} \overline{W}_{s'k}\right)}{\sum_{i' \in M} \exp\ \left(g_{si'} + \sum_{s' \in N_s} \sum_{k=1}^{M} J_{i'k} \overline{W}_{s'k}\right)} \ . \tag{2.36}$$

A derivation of the above equations is presented in Appendix A.

It is worth observing that by eliminating neighborhood interactions from our Gibbs random field specification i.e. by specifying a neighborhood system in which N_s is the empty set for all values of s, we can achieve the update equations of the EM-segmentation method. Thus EM-MF segmentation can be thought of as a generalization of EM-segmentation.

2.3.1 Summary of EM-MF Segmentation

We have described a segmentation method for MRI, EM-MF, that introduces Gibbs modeling of the spatial structure of tissues in conjunction with a Mean-Field solution technique. By Gibbs modeling of the homogeneity of tissue, resistance to thermal noise in the images is obtained. The image data and intensity correction are coupled by an external field to an Ising-like tissue model, and Mean-Field equations are used to obtain posterior estimates of tissue probabilities. The method is a generalization of an existing segmentation framework [Wells III et al.1996], and leads to a modification that is a simple and inexpensive relaxation-like update (Equation 2.36). Unlike other approaches to solving similar problems that have neighbor interactions [Besag1986, Held et al.1998], the present method does not make hard assignments of tissue class, but instead stays within the framework of calculating posterior probabilities on tissue assignments. This approach is consistent with the principle of least commitment, which advocates (when appropriate) the avoidance of information loss that is entailed by thresholding [Marr1982].

In the rest of this chapter, we show the relationship between EM-MF segmentation and the classical Relaxation Labeling technique, and conclude with sample results of using EM-MF segmentation on noisy brain MRI.

2.4 Relaxation Labeling and EM-MF

Relaxation labeling (RL) was introduced in [Rosenfeld, Hummel and Zucker1976] as a mechanism for reducing ambiguity among objects in a scene. It is an approach for minimizing the energy function $E(f)$. In RL the discrete label classification problem is approximated with a continuous label one; sites are allowed to have "soft" (continuous) labels for all classes rather than discrete ("hard") classification into a particular class. It is attractive because of its polynomial time complexity - $O(N^2M^2)$ for a lattice of N sites and M labels.

We use the following notation to describe RL:

- $\mathcal{I}$ is the input image
- N is the number of voxels in $\mathcal{I}$.
- $\mathcal{S} = \{s|1 \leq s \leq N\}$ is a set of integer indices into $\mathcal{I}$.
- M is the number of classes in the segmentation problem.
- U_{sj} is a continuous valued random variable $\in (0, 1)$ that indicates the membership of a voxel s to class j. The magnitude of U_{sj} indicates the strength of the belief that voxel s belongs to class j. A feasibility constraint on the U's dictates that: $\sum_{j=1}^{M} U_{sj} = 1, \forall s \in S$.

- U_s is an M-dimensional random column vector whose j^{th} component is the random variable U_{sj}: $U_s = [U_{s1}, U_{s2}, \ldots, U_{sM}]^T$.
- U is an $N \times M$ dimensional matrix of all continuous label vectors for voxels in the image $\mathcal{I}$: $U = [U_1 U_2 \ldots, U_N]^T$.
- $r_1(s, i)$ is a unary *compatibility function* designed to indicate the compatibility of label i at site s.
- $r_2(s, s', i, i')$ is a binary *compatibility function* designed to indicate the compatibility of labels i and i' at sites s and s' respectively.
- $G(U)$ is the *Gain* of a labeling U. It is defined as the sum of compatibility functions

  $$G(U) \equiv \sum_{s\in\mathcal{S}} \sum_{i=1}^{M} r_1(s,i) U_{si} + \sum_{s\in\mathcal{S}} \sum_{i=1}^{M} \sum_{s'\in\mathcal{S}} \sum_{i'=1}^{M} r_2(s,s',i,i') U_{si} U_{s'i'} \tag{2.37}$$

  and is maximized in the optimization procedure.

The *Relaxation* iteration consists of updating the fuzzy labels U_{si} using gradient ascent on the gain function. Assuming symmetric binary compatibility functions i.e. assuming that $r_2(s, s', i, i') = r_2(s, s', i', i)$. The gradient of the gain function G with respect to the label U_{si}, denoted by q_{si}, can be written as:

$$\begin{aligned} q_{si} &\equiv \frac{dG(U)}{dU_{si}} && (2.38) \\ &= r_1(s,i) + 2 \sum_{s'\in\mathcal{S}} \sum_{i'=1}^{M} r_2(s,s',i,i') U_{s'i'} && (2.39) \end{aligned}$$

The actual gradient ascent is done using an update operator ϕ which is defined as a function of the labeling at the given time step and the gradient. Starting from an initial labeling $U^{(0)}$, it is used to generate successive labelings $U^{(1)}, U_{si}^{(2)}, \ldots$:

$$U^{(t+1)} \leftarrow \phi(U^{(t)}, q^{(t)}) \tag{2.40}$$

Relaxation schemes differ in their choice of the update operator ϕ. For example, Rosenfeld et al. define the update such that:

$$U_{si}^{(t+1)} \leftarrow \frac{U_{si}^{(t)}[1 + q_{si}^{(t)}]}{\sum_{i'=1}^{M} U_{si'}^{(t)}[1 + q_{si'}^{(t)}]} \tag{2.41}$$

Note that the normalization in the update ensures that the labelings obey the feasibility constraint i.e. sum to unity.

2.4.1 Relaxation Interpretation of EM-MF

We now show that by defining the appropriate compatibility functions and the update operator, the Mean-Field iteration obtained in the E-Step (Equation 2.36) can be interpreted as a *relaxation labeling* (RL) iteration. Similar observations are mentioned in [Li1995].

Let us define the unary compatibility function $r_1(s, i)$ to be the external field g_{si} that was originally defined in Equation 2.35:

$$r_1(s, i) \equiv g_{si} \tag{2.42}$$

And let us define the binary compatibility function $r_2(s, s', i, i')$ on a pair of sites s and s' with labels i and i' respectively as:

$$r_2(s, s', i, i') \equiv \begin{cases} J_{ii'} & \text{if } s' \in N_s \\ 0 & \text{otherwise} \end{cases} \tag{2.43}$$

where J is the tissue class interaction matrix that was introduced in Equation 2.28. The binary compatibility function is designed to be symmetric i.e. it is designed so that $r_2(s, s', i, i')$ is the same as $r_2(s, s', i', i)$.

Using these definitions of the compatibility functions, and the definition of the gain function from Equation 2.37, we can write the expression for the gradient q_{si} of the gain as:

$$
\begin{aligned}
q_{si} &= r_1(s,i) + \sum_{s'\in\mathcal{S}} \sum_{i'=1}^{M} r_2(s,s',i,i')U_{s'i'} && (2.44)\\
&= g_{si} + \sum_{s'\in\mathcal{S}} \sum_{i'=1}^{M} J_{ii'}U_{s'i'} && (2.45)
\end{aligned}
$$

The above equations were obtained by substituting in the expression for the gradient from Equation 2.38, the expressions for compatibility functions from Equations 2.42 and 2.43.

Now we design an update rule, similar to the one used in [Faugeras and Berthod1981], such that the labels at time $t+1$ are a normalized exponential of the gradient $q^{(t)}$ at time t:

$$
\begin{aligned}
U_{si} &\leftarrow \frac{\exp\ (q_{si}^{(t)})}{\sum_{k=1}^{M} \exp\ (q_{sk}^{(t)})} && (2.46)\\
&= \frac{\exp\ (g_{si} + \sum_{s'\in\mathcal{S}} \sum_{k=1}^{M} J_{ik}U_{s'k})}{\sum_{i'=1}^{M} \exp\ (g_{si'} + \sum_{s'\in\mathcal{S}} \sum_{k=1}^{M} J_{i'k}U_{s'k})} && (2.47)
\end{aligned}
$$

Note that this is the same as the update Mean-Field update equation that we proposed in Equation 2.36, with the change of variables from U_{si} to $\overline{W}_{si}$. Thus, we have shown here that the update rule used in the proposed EM-MF segmentation algorithm (Equation 2.36) is formally equivalent to a *relaxation labeling* update of Equation 2.46.

2.5 Experiments

In an experiment designed to demonstrate the robustness of EM-MF segmentation we prepare an image by adding a substantial amount of Gaussian i.i.d. noise to a gradient echo MRI scan. This additive noise simulates an exaggeration of the thermal noise present in MRI scans. The results of EM-MF segmentation on this noisy image are shown in Figure 1. Note how the EM-MF approach effectively suppresses the effects of this noise in the segmentation while preserving structural details.

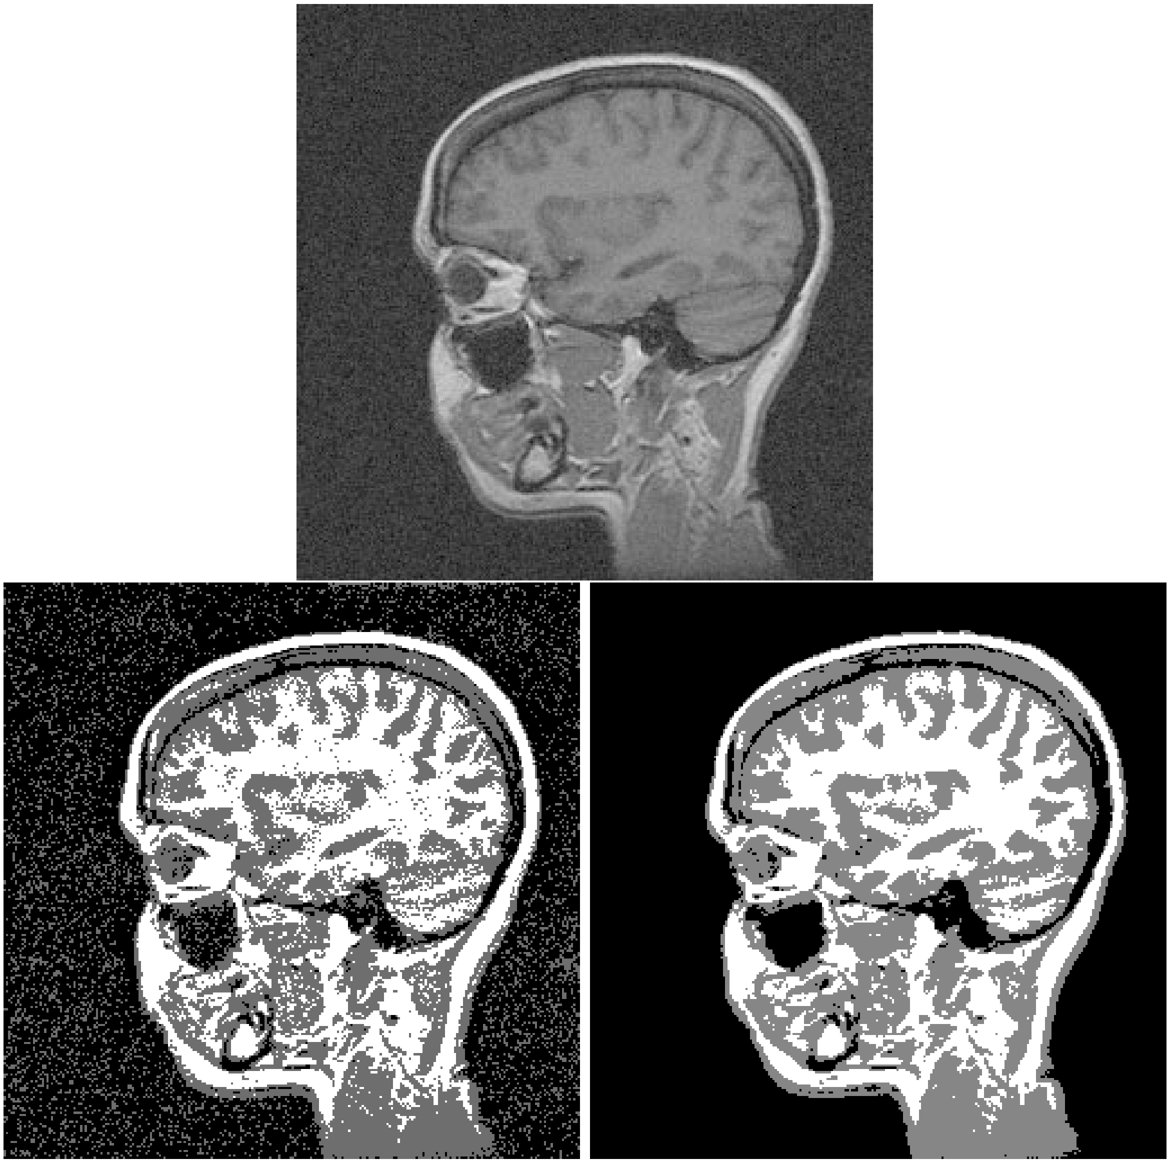

Figure 2-8: This experiment has been designed to demonstrate the robustness of EM-MF segmentation with respect to thermal noise in MRI. The top image was prepared by adding a substantial amount of Gaussian i.i.d. noise to a gradient echo MRI scan. The bottom left image shows the result of standard EM-segmentation (3d) into three classes: white matter, gray matter, and background. The bottom right image shows the EM-MF segmentation (3d). In both segmented images, white matter is shown in white, gray matter is shown in gray, and the background/air is shown in black. Note how the EM-MF approach has effectively suppressed the effects of this noise in the segmentation while preserving structural details.

Chapter 3

Modeling Geometry of Structures

In this chapter we introduce *geometric models* for anatomical structures. These models provide a formalism for encoding notions of global spatial relationships such as "articular cartilage lines the bones" between bones and cartilage. *Geometric models* are intended to be used in segmentation of structures that are easiest to distinguish in an image based on their spatial relationship to other structures.

As an example, consider the artistic rendering of knee anatomy shown in Figure 3-1. Structures such as the articular cartilage, the ligaments and the tendons are well described in terms of their relationship to the bones. Now examine the articular femoral cartilage in the cross-sectional slice of a T2-weighted knee MRI scan shown in Figure 3-2. It is a small structure that occupies less than one hundredth of the total knee volume in a typical scan, and its graylevel appearance is very similar to more than thirty percent of the rest of the image. However, radiologists and orthopedic surgeons have little trouble identifying it in these images. And they certainly never mistakenly identify muscle, or trabecular bone – both of which have similar graylevel patterns – as cartilage. This is simply because they are experts in the layout of the knee, and specifically know the fact that articular cartilage lines the bones. It is this level of information about the spatial relationships of structures that we seek to encode via *geometric models*.

Before proceeding with details of the proposed *geometric models*, it is worthwhile to examine some other ways in which geometric information can be used by an identi-

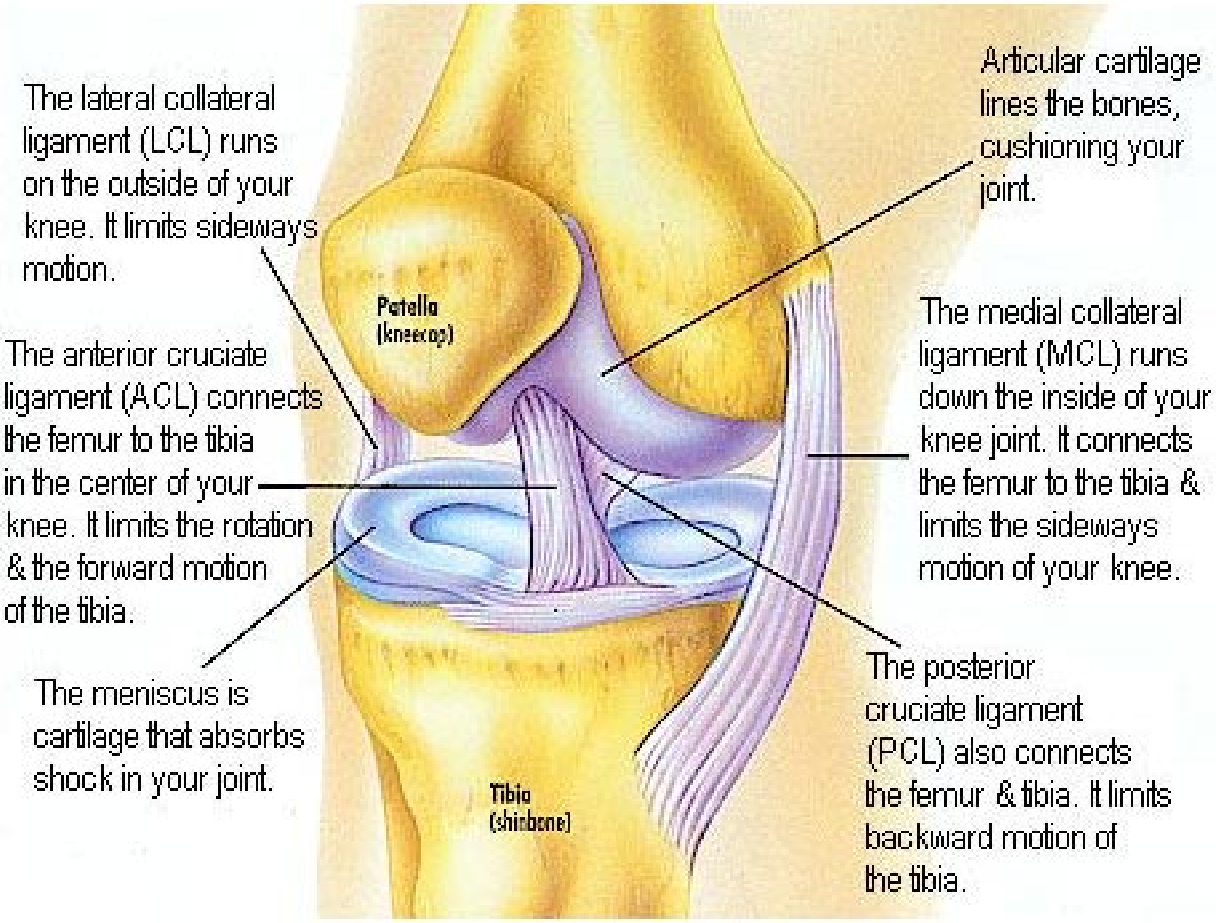

Figure 3-1: The anatomy of a knee [Hinton].

fication program. *Connectivity-analysis*, in which particular four- or eight-connected components in the image are assigned anatomically meaningful labels, is a common method for encoding topological information in the segmentation process (see textbooks [Horn1986, Pavlidis1982] for background material). Morphological operations such as *opening* and *closing* are used to isolate convex structures by shaving off thin connectors emanating from them (see textbook [Serra1982] for background material). Deformable models derived from *snakes* [Kass, Witkin and Terzopoulos1988] search for the surface of a structure in the vicinity of a given initial configuration. In short, there are well-known techniques that can be used to implicitly supply geometric knowledge to a segmentation program on a case-by-case basis.

In the rest of this chapter, we describe the concept of *geometric models* (Section

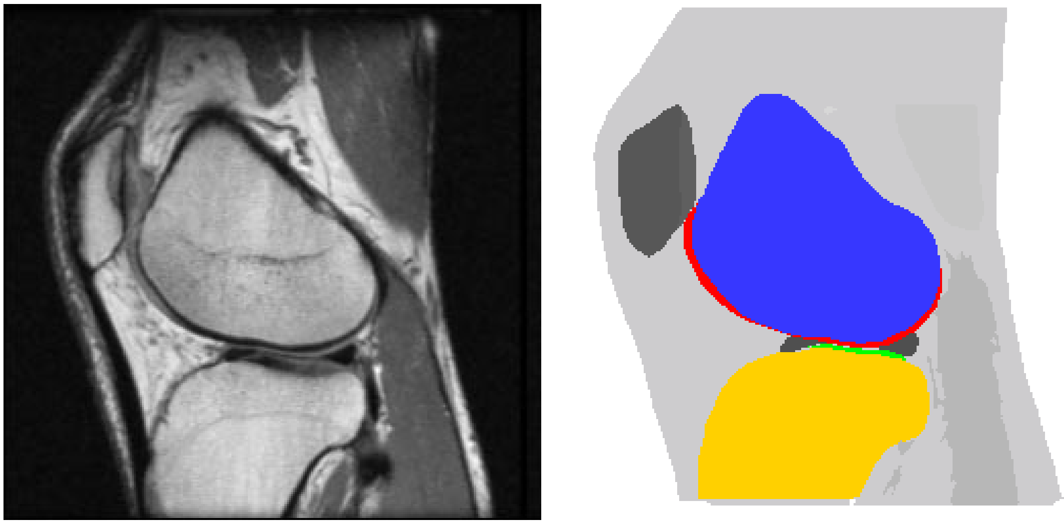

Figure 3-2: On the left is a sagittal cross-section image from a knee MRI. On the right is a segmentation of the image into four structures: the femur - the thigh bone (blue), the tibia - the shin bone (yellow), the articular femoral cartilage (red), and the tibial cartilage (green). Note the undistinguishing graylevel of the cartilage in the image on the left, and its unique geometric relationship to the bone in the image on the right.

3.1); instantiate the concept by presenting two specific geometric models (Section 3.2 and 3.3); mention the role of such models in statistical classification (Section 3.4), and discuss heuristics for using these models in segmentation (Section 3.5).

3.1 Defining Geometric Models

We represent a three-dimensional image $\mathcal{I}$ with R rows, C columns and D slices as a point set of coordinate locations:

$$\mathcal{I} = \{s = (x_s, y_s, z_s) | 1 \leq x_s \leq C, 1 \leq y_s \leq R, 1 \leq z_s \leq D\}, \tag{3.1}$$

and we use "voxel s" to refer to the coordinate location s in $\mathcal{I}$.

We define a *geometric model* for an anatomical structure as a predictive probability distribution on coordinate locations in $\mathcal{I}$. This model predicts the probability of membership of each coordinate location in an image, conditioned on the coordinate locations of one or more other structures in the image.

In order to further develop the notion of a geometric model, we define some notation and terminology:

- A *structure* is an unordered subset of coordinate locations of $\mathcal{I}$. "Structure j" refers to all coordinate locations s with label j in a labeled image.
- *Primary structures* in $\mathcal{I}$ are those structures upon whose segmentation geometric models for *secondary structures* are conditioned. Primary structures need to be segmented from $\mathcal{I}$ before geometric models can be constructed for secondary structures.
- *Secondary structures* are those structures in $\mathcal{I}$ for which geometric models are to be created by conditioning on segmentation of primary structures.
- M' is the number of primary structures in $\mathcal{I}$.
- L' is the set $\{l'_1, l'_2, \ldots, l'_{M'}\}$ of labels of the M' primary structures in $\mathcal{I}$.
- M is the number of secondary structures in $\mathcal{I}$.
- L is the set $\{l_1, l_2, \ldots, l_M\}$ of labels of the M secondary structures in $\mathcal{I}$.
- N is the number of voxels in $\mathcal{I}$.
- W' is a segmentation of $\mathcal{I}$ into the mentioned M' primary structures. It is an $N \times M'$ matrix whose sj^{th} element W_{sj} is a binary indicator random variable that indicates whether voxel s belongs to class l'_j or not:

  $$W'_{sj} = \begin{cases} 1 & \text{if voxel } s \text{ belongs to class } l'_j \\ 0 & \text{otherwise} \end{cases}$$

  The state space of W'_s is the canonical basis for $\mathcal{R}^{M'}$.
- W is a segmentation of $\mathcal{I}$ into two structures: l and $\bar{l}$. It is a $N \times 2$ matrix in which element W_{s1} is a binary indicator variable that indicates whether voxel s

belongs to class l or not:

$$W_{s1} = \begin{cases} 1 & \text{if voxel } s \text{ belongs to class } l \\ 0 & \text{otherwise} \end{cases}$$

Similarly, element W_{s2} is a binary indicator variable that indicates whether voxel s belongs to class $\overline{l}$ or not, and is the logical *not* of W_{s1}.

Using the notation and terminology defined above, we state the definition of a geometric model more precisely as follows:

Given a coordinate location $s \in \mathcal{I}$, and a labeling W' of $\mathcal{I}$ into primary structures L', a *geometric model* for structure l is a probability distribution function

$$P(s \in l|W') \tag{3.2}$$

that predicts for each voxel $s \in \mathcal{I}$ the probability that it belongs to structure l, using the coordinate location of s and the coordinate locations of the previously segmented primary structures in $\mathcal{I}$.

Two characteristics of geometric models are worth making explicit at this point to help the reader understand their applicability to segmentation:

- *Primary Structures are required.* Geometric models are created using segmentation of primary structures. They can not be used if segmented primary structures are not available for bootstrapping the segmentation.

- *Selection of Primary and Secondary structures is important.* In order to successfully create and apply geometric models in segmentation, one needs to understand the attributes that are desirable when selecting primary and secondary structures. As a heuristic, structures chosen to be primary should be easy to identify in images using low-level image processing techniques. Secondary structures should be such that an anatomist would describe them to a lay person in terms of their spatial relationships to primary structures. In subsequent chapters we will briefly discuss methods for automating the optimal selection of these structures. However, in this work the structures will be selected manually.

In the next two sections we propose two geometric models: a *distance model* (Section 3.2) and a *distance-normal model* (Section 3.3) for encoding spatial relationships between structures. For each of these geometric models, we present in the respective sections, a description of the particular model, an algorithm for estimating the probability distribution from training images, and an example to illustrate the density estimation process.

3.2 Distance Model

A *distance model* predicts the probability of the event that voxel s belongs to secondary structure l, given the Euclidean distance between the coordinate location of s and the primary structures in the image. We define it as a conditional probability distribution on coordinate locations s in l:

$$P(s \in l | \rho_{s,l'_1}, \rho_{s,l'_2}, \ldots, \rho_{s,l'_M}). \tag{3.3}$$

Here, ρ_{s,l'_j} denotes the distance from voxel s to primary structure l'_j, and is defined as the L^2 norm of the line segment connecting s to the closest point in structure l'_j. We illustrate the concept of distances between voxels and structures in Figure 3-3. The figure shows example structures l_1 and l_2 as well as distances between them. Structure l_1 is the set of points marked with circles, and l_2 is the set of points marked with stars. The length of the line segment connecting the s^{th} point in l_2 to the closest point in l_1 is its distance ρ_{s,l_1} to l_1.

Note that the PDF for this distance model (Equation 3.3) is derived from the general PDF for the geometric model (Equation 3.2) by replacing the conditioning on the labeling W' with conditioning on distances to primary structures. This reflects our assumption that the distances between structures are a *sufficient statistic* of their relative geometry, meaning that if we know the distances between the secondary

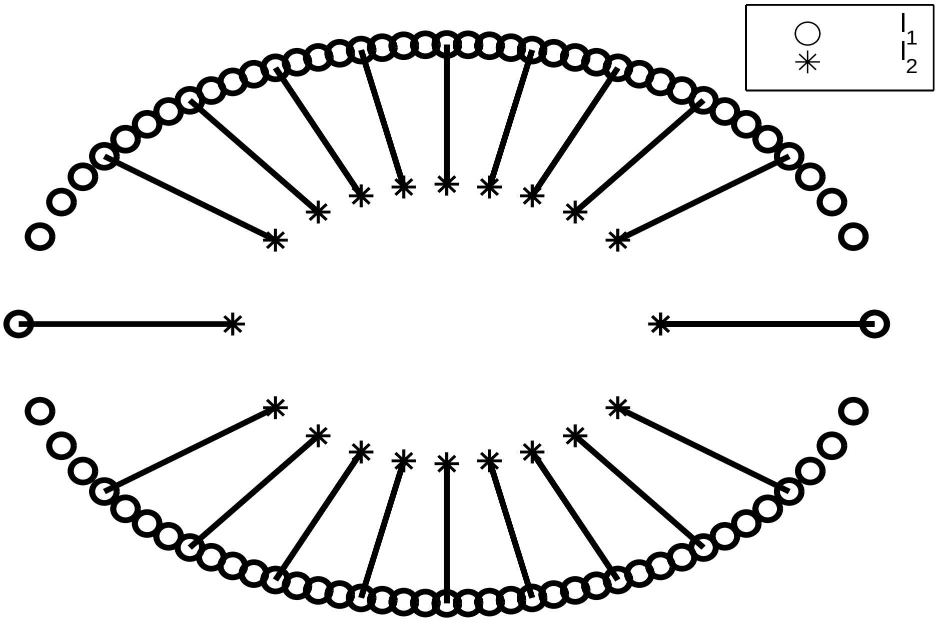

Figure 3-3: Two structures l_1 and l_2 are shown here. l_1 consists of the points marked with circles, and l_2 consists of the points marked with stars. The distance between a point on l_2 and the structure l_1 is the length of the line segment, as shown here, connecting that point to the closest point in l_1.

structure and the primary structures, we assume that we know everything about their spatial relationship that is necessary to the problem at hand. It is important to remember this assumption when deciding whether or not to apply this distance model for a segmentation problem.

3.2.1 Algorithm for Estimating a Distance Model

The distance model of Equation 3.3 is estimated from labeled images. We first describe how to estimate a distance model from a single, labeled, three-dimensional image. Then we show how to extend the estimation procedure to include multiple labeled images.

To estimate a distance model for secondary structure l, we use as input a single image in which primary structures $l'_1, l'_2, \ldots l'_{M'}$, and secondary structure l have been labeled. To estimate the distance model, we

1. Compute the distances ρ_{s,l'_j} for all voxels s in the secondary structure l to each of the primary structures l'_j in L'. The distances computed are signed so that all points outside the primary structures get positive distance values, while the points inside get negative distance values.

2. Estimate the M'-dimensional class conditional density function

$$P(\rho_{s,l'_1}, \rho_{s,l'_2}, \ldots \rho_{s,l'_{M'}} | s \in l)$$

   on distances $\rho_{s,l'_1}, \rho_{s,l'_2}, \ldots \rho_{s,l'_{M'}}$ (from secondary structure l to each of the primary structures) using the samples of ρ_{s,l'_j} computed in Step 1. This PDF can be estimated from the samples as a parametric (for example, Gaussian distribution) or non-parametric (for example, Parzen distribution, histogram) form (see textbook [Duda and Hart1973] for density estimation techniques).

3. Use Bayes' rule to compute the distance model PDF from the class conditional density:

$$P(s \in l | \rho_{s,l'_1}, \rho_{s,l'_2}, \ldots, \rho_{s,l'_M}) = \frac{P(\rho_{s,l'_1}, \rho_{s,l'_2}, \ldots, \rho_{s,l'_M} | s \in l) P(s \in l)}{\sum_{l' \in L} P(\rho_{s,l'_1}, \rho_{s,l'_2}, \ldots, \rho_{s,l'_M} | s \in l') P(s \in l')} \tag{3.4}$$

   Here the expression on the left hand side is the distance model from Equation 3.3. In the right hand side expression, the first term in the numerator is the class conditional PDF estimated in Step 2. The second term, $P(s \in l)$, is the prior probability that a voxel belongs to structure l . This could be the spatially stationary priors that were used in Chapter 2. The denominator acts as a normalization factor.

To estimate a distance model for secondary structure l from multiple training images, we normalize the distances from each image such that the minimum and maximum distance values are the same. All the normalized distances are used to build the probability distribution function.

3.2.2 Example of Estimating a Distance Model

Let us instantiate the model estimation procedure from a single labeled image in the simple case shown in Figure 3-4 of a two-dimensional image with two structures: l' is a circle of radius 5 centered at the center of the image, and l is a larger circle of radius 10 with the same center (both structures consist of 1000 points each). Here the smaller circle l' is the primary structure in the image that is segmented in the training image and will be segmented in all novel images. The larger circle l is the secondary structure that is segmented only in the training image, and it is for l that we will estimate a distance model which will be used to segment it in novel images. For the purpose of this example, we will assume that a novel image in which l' is already segmented, is to be classified into two classes: the secondary structure l and $\bar{l}$. To estimate the distance model for l, we follow the steps outlined in the previous section:

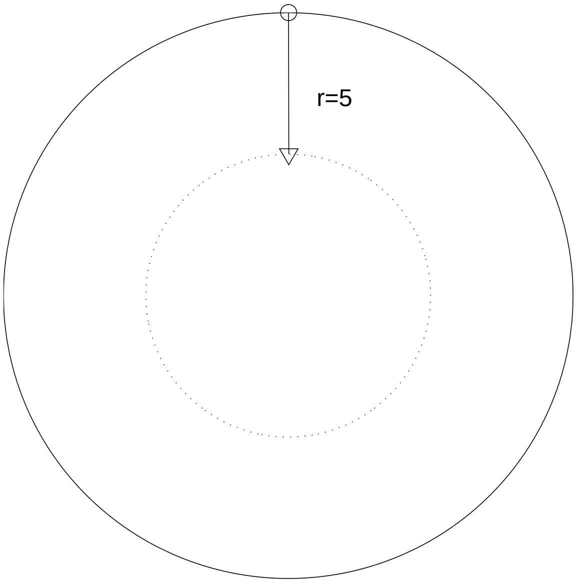

Figure 3-4: An example to illustrate the estimation of a distance model. This example image shows a primary structure l' - the larger circle, and a secondary structure l - the smaller circle. The goal is to estimate a distance model for the secondary structure, given the segmentation of the primary structure.

1. The distance is 5 for all points in l, so $\rho_{s,l'} = 5, \qquad \forall s \in l$.

2. We estimate the class conditional density on distances for structure l as a normalized histogram of distances $\rho_{s,l'}$. The histogram is a bar of height 1000 at $\rho_{s,l'} = 5$ (as shown in Figure 3-5) and the estimated class conditional density is the Dirac delta function $\delta(\rho_{s,l'} - 5)$.

3. Note that this is a two-class problem and therefore the conditional density on $\rho_{s,l'}$ for the class $\bar{l}$ is uniform on the intervals $[0, 5)$ and $(5, 70]$ (70 is the maximum possible distance in this case) and is zero at $\rho_{s,l'} = 5$ (as shown in Figure 3-6). We use Bayes' rule to get the distance model as follows:

$$\begin{aligned} P(s \in l|\rho_{s,l'_1}) &= \frac{P(\rho_{s,l'_1}|s \in l)P(s \in l)}{P(\rho_{s,l'_1}|s \in l)P(s \in l) + P(\rho_{s,l'_1}|s \in \bar{l})P(s \in \bar{l})} \qquad (3.5) \\ &= \delta(\rho - 5) \end{aligned}$$

which is the Dirac delta at $\rho = 5$.

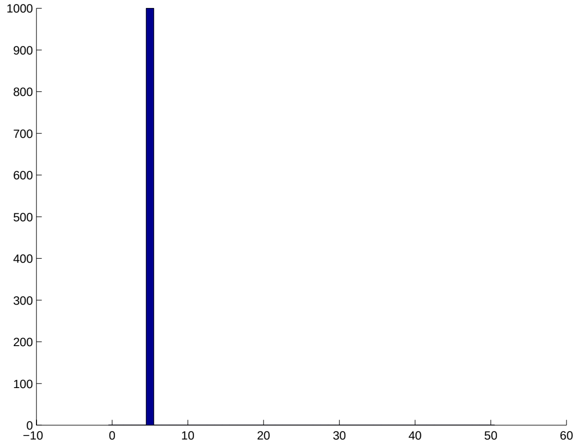

Figure 3-5: A histogram of distances from points in secondary structure l to primary structure l'. The estimated class conditional probability density function for l resulting from normalizing the histogram is the Dirac delta at 5.

The expression for the distance model (Equation 3.5) has the straightforward interpretation that the secondary structure l appears at a distance 5 from l'. This example illustrates the fact that relative geometry of structures, such as the ones in

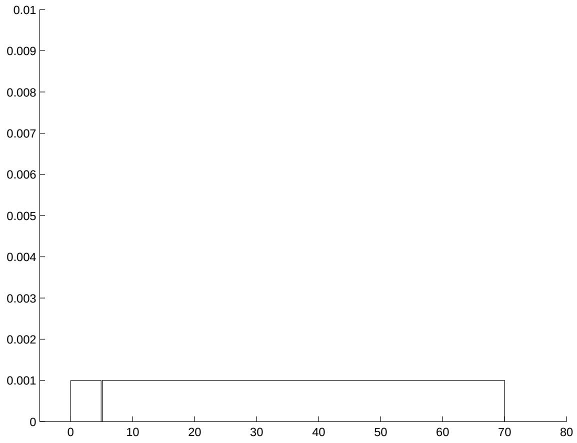

Figure 3-6: The estimated class conditional density function for $\bar{l}$. Note that it is uniform on the intervals $[0, 5)$ and $(5, 70]$ and is zero at $\rho_{s,l'} = 5$.

Figure 3-4, can be approximated in a compact manner using this distance model. We say "approximated" here for multiple reasons. First, the density is estimated from samples, hence is an estimate of the actual density. Second, the distance model may not capture the exact relative geometry of the two structures, but may capture enough to still be useful for segmentation purposes.

3.2.3 Summary of Distance Model

In summary, the proposed distance model captures a distance based geometric relationship between primary and secondary structures. In this section we described this model, presented an algorithm for estimating it from labeled training images, and instantiated the model estimation algorithm for a particular example in which the primary and secondary structures were two concentric circles.

In the next section we augment our measurement of the relationship between the primary and secondary structures using distances as well as local surface orientations. We call the resulting model a *distance-normal model* for geometric relationships.

3.3 Distance-Normal Model

Our *distance-normal* model takes into account not only the $\mathcal{L}^2$ norms between points in the secondary structure l and the closest points in the primary structure l', but it also factors in the direction of the local normal vectors to the surface of l'. We define it as a conditional probability distribution on coordinate locations s in l:

$$P(s \in l | \rho_{s,l'_1}, n_{s,l'_1}, \rho_{s,l'_2}, n_{s,l'_2}, \ldots, \rho_{s,l'_{M'}}, n_{s,l'_{M'}}). \tag{3.6}$$

As was the case in the distance model, ρ_{s,l'_j} here denotes the signed distance from voxel s to the closest point in the primary structure l'_j. We refer to this closest point as q_{s,l'_j}. n_{s,l'_j} denotes the local normal to the surface of the primary structure l' at q_{s,l'_j}. We illustrate this concept of distances and normals relating the geometry of two structures in Figure 3-7. The figure shows example structures l_1 (a circle) and l_2 (an arc lying below the circle) and the distances and normals relating their geometry.

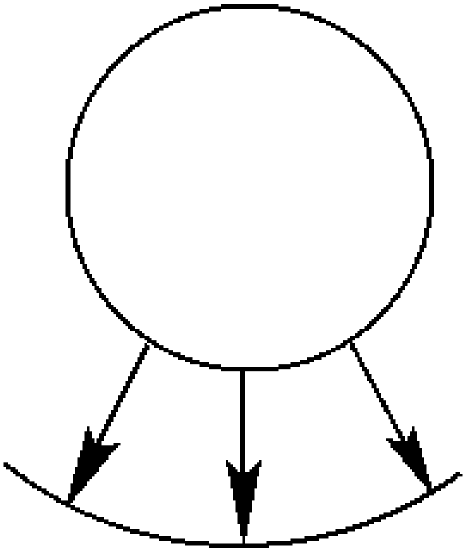

Figure 3-7: Two structures l_1 and l_2 are shown here. l_1 is the circle and l_2 is the arc lying below it. Line segments are shown between some points in l_2 and the closest points on l_1. The direction of the local surface normals to l_1 is also shown.

The PDF for our distance-normal model (Equation 3.6) is derived from the general PDF for the geometric model (Equation 3.2) by replacing the conditioning on the labeling W' with conditioning on distances to primary structures as well as the local normals of the primary structures. This reflects our assumption that in addition to the distances between two structures, the orientations of local surface patches of the primary structure offer valuable predictive power for the geometry of secondary structures. Intuitively, if the distance model can be used to describe geometric relationships such as “structure S lies about 5-10mm inside/outside structure P”, then the distance-normal model is better suited for relationships such as “structure S lies 5-10mm outside of and above structure P”.

3.3.1 Algorithm for Estimating a Distance-Normal Model

We first describe how to estimate a distance-normal model from a single, labeled, three-dimensional image. Then we extend the estimation procedure to include multiple labeled images.

To estimate a distance-normal model for secondary structure l, we use as input a single image in which primary structures $l'_1, l'_2, \ldots l'_{M'}$, and secondary structure l have been labeled. To estimate the distance-normal model, we:

1. Compute the signed distances ρ_{s,l'_j} for all voxels s in the secondary structure l to the closest points, $q_{s,l'j}$, on each of the primary structures l'_j in L'.

2. Compute the normal n_{s,l'_j} to the primary structures at the closest points q_{s,l'_j}.

3. Estimate the $3M'$-dimensional[1] class conditional density function

$$P(\rho_{s,l'_1}, n_{s,l'_1}, \rho_{s,l'_2}, n_{s,l'_2}, \ldots, \rho_{s,l'_M}, n_{s,l'_M} | s \in l)$$

   using the samples of distances ρ_{s,l'_j} and normals n_{s,l'_j} computed in Step 1. As was the case with the distance model, this PDF can be estimated from the samples as a parametric or a non-parametric form.

---

[1] For three dimensional data, a normal direction is expressed using two parameters, and the third parameter is for the distance.

4. Use Bayes' rule to compute the distance-normal model PDF from the class conditional density:

$$P(s \in l | \rho_{s,l'_1}, n_{s,l'_1}, \ldots, \rho_{s,l'_{M'}}, n_{s,l'_{M'}}) = \frac{P(\rho_{s,l'_1}, n_{s,l'_1}, \ldots, \rho_{s,l'_{M'}}, n_{s,l'_{M'}} | s \in l) P(s \in l)}{\sum_{l' \in \{l, \bar{l}\}} P(\rho_{s,l'_1}, n_{s,l'_1}, \ldots, \rho_{s,l'_{M'}}, n_{s,l'_{M'}} | s \in l') P(s \in l')} \tag{3.7}$$

   The left hand side is the distance-normal model from Equation 3.6. On the right hand side, the first term in the numerator is the class conditional PDF estimated in Step 2. The second term, $P(s \in l)$, is the prior probability that a voxel belongs to structure l. The denominator is the usual normalizer to ensure that the numbers on the right hand size add up to unity.

To estimate a distance-normal model for secondary structure l from multiple training images, two normalization steps are required. The distances are normalized across images by scaling them such that they lie within the same range. Normalization of the orientations, however, can be done in one of several ways. One approach is to normalize for the orientation of one primary structure by aligning the principal axes of the primary structure across images. Another approach is to assume that the primary structures vary in a non-rigid manner across images, and to align each one independently.

3.3.2 Example of Estimating a Distance-Normal Model

Let us instantiate the model estimation procedure from a single labeled image in the case shown in Figure 3-8 of a two-dimensional image with two structures: the primary structure l' is a circle of radius 5 centered at the center of the image, and the secondary structure l is a concentric arc of radius 10, lying below it. We now estimate a distance-normal model for the arc l relative to the primary structure circle. As before, we assume that a novel image is to be classified into two classes: l and $\bar{l}$.

1. The distance is 5 for all points in l, so $\rho_{s,l'} = 5, \qquad \forall s \in l$.

2. This example is in two-dimensions and the normal direction is expressed using

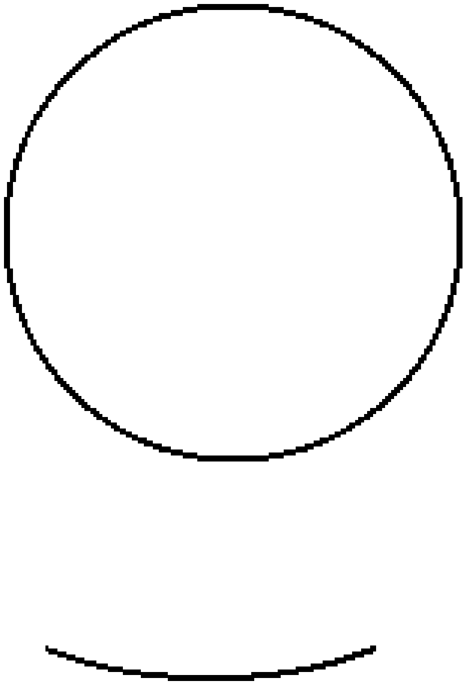

Figure 3-8: An example to illustrate the estimation of a distance model. This example image shows a primary structure l' - the circle, and a secondary structure l - the arc lying below the circle. The goal is to estimate a distance-normal model for the arc, given the location of the circle.

a single parameter, the angle the normal vector makes with the x-axis. The normals range from $\pi/6$ to $5\pi/6$, so $\pi/6 \leq n_{s,l'} \leq 5\pi/6 \qquad \forall s \in l$.

3. The class conditional density on distances and normals is a two-dimensional distribution. This is because the distance and normals contribute one dimension each, and the number of primary structures, M', is 1 in this example. We estimate this density for structure l using a normalized histogram of the samples of distances and corresponding local normal directions that were computed in the previous step. This density is uniform for $\rho_{s,l'} = 5$ and $\pi/6 \leq n_{s,l'} \leq 5\pi/6$.

4. Note that this is a two-class problem and therefore the conditional density on for the distances and normals for the class $\bar{l}$ is the flip of the PDF for l (i.e. 1 minus PDF for l). We use Bayes' rule to get the distance-normal model as follows:

$$P(s \in l|\rho_{s,l'}, n_{s,l'}) = \frac{P(\rho_{s,l'}, n_{s,l'}|s \in l)P(s \in l)}{P(\rho_{s,l'}, n_{s,l'}|s \in l)P(s \in l) + P(\rho_{s,l'}, n_{s,l'}|s \in \bar{l})P(s \in \bar{l})} \tag{3.8}$$
$$= \begin{cases} \frac{3}{2\pi}\delta(\rho - 5), & \pi/6 \leq n_{s,l'} \leq 5\pi/6 \\ 0 & \text{otherwise} \end{cases}$$

The above expression for the distance-normal model is a compact representation of the relationship that is approximately described in words as "the arc lies 5 voxels below the circle".

3.3.3 Summary of Distance-Normal Model

In summary, the proposed distance-normal model augments the distance model that we presented earlier by using additional information about local surface normals to relate the geometry of structures. In this section we have described this model, presented an algorithm for estimating it from labeled training images, and instantiated the model estimation algorithm for a particular example in which the primary structure was a circle and the secondary structure was a concentric arc lying below it.

3.4 Incorporating Geometric Models in Statistical Classification

Geometric models, such as the distance model (Equation 3.3) and the distance-normal model (Equation 3.6) are intended to be used as spatially varying prior probability distributions on secondary structures in Bayesian classification algorithms. In the next chapter, we will show two specific applications in which we use geometric models in statistical classification. In one application, the statistical classifier used is a standard Bayesian MAP classifier. In the other application we incorporate a geometric model in the Expectation-Maximization based EM-MF classifier that we presented in Chapter 2.

3.5 Heuristics for Judging Efficacy of a Geometric Model

The appropriateness of a given geometric model for a particular task may be determined using two heuristics: the information content, in the information theoretic sense, of the geometric model, and the discrimination power of the model(see textbooks [Cover and Thomas1991, Duda and Hart1973] for background in Information Theory and Discriminant Analysis). We discuss each of these heuristics next.

3.5.1 High Information/ Low Entropy Geometric Model

Geometric relationships with high information content are the ones that are best exploited in segmentation using geometric models. The PDFs for such relationships have low entropy i.e. they are "peaky" as opposed to uniform. A high entropy, or low information, distribution for the geometric model indicates the lack of a measurable relationship, across a population, between the structures. Such a scenario arises when either there really isn't any significant underlying geometric relationship between the structures, or that the particular geometric model being used is inadequate. Either case is reason to re-evaluate the geometric model that is being applied.

3.5.2 High Discrimination Power

Recall that the end goal of the segmentation problem, for which geometric models are used, is to discriminate between different structures. Geometric models, thus, need to have maximal power to discriminate between these structures. For example, distance models are most effective for segmentation when they divide the distance-space into non-overlapping regions for each of the secondary structures. Recall that since the distance model is used as a prior distribution on a secondary structure, the ability to rule out entire regions of the image based simply on distance to a primary structure can be very useful to an otherwise intensity based (geometry-free) classification mechanism. In the next chapter we illustrate the utility of a distance model in segmentation of brain tissue from MRI images in conjunction with the EM-MF method of Chapter 2.

No matter how low the entropy, distance models are not much use in segmentation of structures that they can't discriminate between. Two secondary structures can have very different underlying geometry, and yet generate the same distance model. It is for these structures that the distance model offers no discriminating power and should not be used. This point is illustrated in the configuration of structures in Figure 3-9 where the primary structure is a circle of radius 5, and the first secondary structure l_1 is a concentric arc of radius 10 lying above the circle, and the second secondary structure is a concentric arc, also of radius 10, lying below the circle. The distance model for each secondary structure is a Dirac delta at 5, hence offering no discrimination power at all.

While using any the geometric models proposed in the chapter we need to keep in mind that these models are many-to-one mappings and the set of different geometries that map into the same geometric model is fairly large.

Continuing with the example image introduced in Figure 3-9, it is useful to note that it makes sense to use a distance-normal model in this case because the directions of the local normals of the primary structure help discriminate between the two secondary structures. A distance-normal model is best used in general when the secondary structures to be distinguished in the end-goal segmentation divide the distance-normal-space into non-overlapping regions. In the next chapter we will illustrate the utility of a distance-normal model for segmentation of articular cartilage from knee MRI.

It is also worth pointing out here that the distance-normal model doesn't help much when the *mutual information* between the distances and the normals is low, e.g. the case when the normals are uniformly distributed per distance value. If the secondary structure is in the form of rings around a primary structure, no information is offered by the normals in addition to the distances, and there is no advantage to using the distance-normal model.

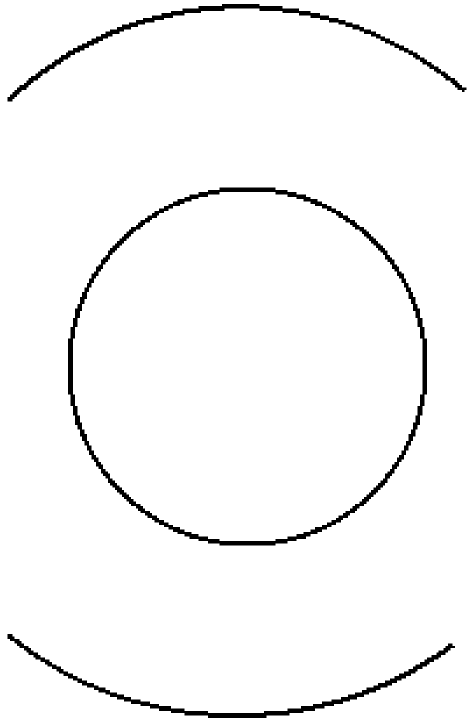

Figure 3-9: . An example to illustrate the inappropriateness of a distance model. The circle of radius 5 is the primary structure in this case, and the concentric arc of radius 10 lying below the circle is one secondary structure and the concentric arc also of radius 10 but lying above the circle is the other secondary structure. It turns out that the distance model in each case is a Dirac delta at 5, and hence offers no discrimination power between the two structures.

3.6 Summary of Geometric Models

In this chapter we introduced *Geometric Models* which are a statistical encoding for the relative geometry of structures across a population. Two specific instantiations of this concept of geometric models are: the distance model, and the distance-normal model. These models are useful for capturing spatial relationships that are commonly used to describe anatomical structures relative to each other: “structure A lies near structure B”, “A lies above (below, to the left of, to the right of, in front of, to the back of) B”, “A lies about 10 mm below B”. In order to optimally use a particular geometric model in a given application, one needs to be keep in mind the large equivalence class of underlying geometries that are mapped into the same geometric model. In the next chapter we illustrate the utility of geometric models to the segmentation of medical imagery. Specifically, we will use the distance model in the segmentation of brain MRI, and a distance-normal model in the segmentation of knee MRI.

Chapter 4

Experiments and Results

In this chapter we report experiments in segmentation of MR images that illustrate the performance of the two contributions of this thesis: the EM-MF segmentation method that was described in Chapter 2, and the Geometric Models that were described in Chapter 3. The first experiment is designed to demonstrate the performance of the EM-MF segmentation algorithm in de-noising of brain MRI. The second experiment is designed to illustrate the construction of Geometric Models for structures in the brain. It also demonstrates the ease with which these geometric models can be used as priors in the EM-MF algorithm to produce an integrated segmentation method that uses knowledge about the gray-level appearance of tissues, accounts for the inhomogeneity and thermal noise in MRI, and uses information about the relative geometry of structures in the human head. The third experiments illustrates the generalizability of Geometric Models and demonstrates their use in segmentation of knee MRI.

4.1 EM-MF Segmentation for Brain MRI

In Chapter 2 we described the EM-MF segmentation method for MR images using Gibbs modeling of tissues in conjunction with a Mean-Field solution technique. By Gibbs modeling of the homogeneity of tissue, we obtain resistance to thermal noise in the images. In this section we present an experiment to illustrate the performance of EM-MF segmentation method on noisy brain MRI.

4.1.1 Motivating Applications

The goal of this experiment is to perform an intensity-based classification of a noisy MRI scan into four classes: white matter, gray matter, scalp, and cerebrospinal fluid (CSF). Such segmentations have applications in image-guided surgery and Neuroscience. In image-guided surgery, 3D visualization such as the ones shown in Figure 1-3 in Chapter 1 are created by rendering 3D models of the brain with overlays of tumors for visualization in pre- and intra-operative surgical planning. In Neuroscience, studies are conducted that use segmented images to measure the volumes of white and gray matter in different parts of the brain in order to develop hypotheses about various disorders [Hirayasu et al.1998].

4.1.2 The Input Data

The input scan consists of Gradient Echo MRI with $256 \times 256 \times 128$ voxels, and the voxel size is $1mm \times 1mm \times 0.95mm$ (width×height×depth).

We note that in gradient echo images the intensity appearance of CSF, air and bone are very similar, hence the class "CSF" in this experiment refers to these three underlying tissue classes collectively.

4.1.3 The Experiment

In this experiment we prepared images by adding increasing amounts of Gaussian i.i.d. (independent, identically distributed) noise to input MRI scans. This additive noise simulates an exaggeration of the thermal noise present in MRI scans. We created one set of contaminated images using additive Gaussian i.i.d. noise of standard deviation 5% of the range of intensities in the input image. We refer to these images as "moderately contaminated" images in the rest of this experiment. We created a second set of images by doubling the standard deviation of the Gaussian i.i.d. noise to 10% of the range intensities in the image. We refer to these images as "severely contaminated" images in the rest of this experiment.

4.1.4 Details of EM-MF Implementation for Brain MRI Segmentation

Recall from the description of EM-MF segmentation that the E-Step of the algorithm consists of iterating to convergence Equation 2.36, which is repeated here:

$$\overline{w}_{si} \leftarrow \frac{exp\ (\sum_{s' \in N_s} \sum_{k=1}^{M} J_{ik} \overline{w}_{s'k} + g_{si})}{\sum_{i' \in M} exp\ (\sum_{s' \in N_s} \sum_{k=1}^{M} J_{i'k} \overline{w}_{s'k} + g_{si'})} \tag{4.1}$$

where

- J is the tissue class interaction matrix J, computed from training data as the following relative frequency (repeated from Equation 2.28):

$$J_{kl} = log\ \frac{\#\text{ times class } k \text{ occurs next to class } l \text{ in training data}}{\#\text{ times class } k \text{ occurs in training data}}, \tag{4.2}$$

- g is the external field (repeated from Equation 2.35):

$$g_{si} = f_i + \ log\ P(Y_s | W_s = e_i, \beta_s^{(t)}), \tag{4.3}$$

- f is the logarithm of stationary priors on the tissue class (repeated from Equation 2.29):

$$\begin{aligned} f_i &= log\ p_i^{stat} & (4.4) \\ &= log\ \frac{\#\text{ times class } i \text{ occurs in training data}}{\#\text{ voxels in training data}}, & (4.5) \end{aligned}$$

- $P(Y_s | W_s = e_i, \beta_s^{(t)})$ is the Gaussian measurement model (from Equation 2.14).

$$P(Y_s|W_s = e_i, \beta_s^{(t)}) = \frac{1}{\sqrt{2\pi\sigma_i^2}}\; exp\;(-\frac{(Y_s - \beta_s^{(t)} - \mu_i)^2}{2\sigma_i^2}) \tag{4.6}$$

In the experiments reported in this section, we implemented the above iteration using the following:

- The tissue interaction matrix J was a 4×4 matrix because there are four classes in the classification problem: white matter, gray matter, CSF, and scalp. The value for J that we used is:

  $$J = \begin{bmatrix} 0.90 & 0.07 & 0.03 & 0.01 \\ 0.07 & 0.87 & 0.04 & 0.02 \\ 0.03 & 0.04 & 0.91 & 0.02 \\ 0.01 & 0.02 & 0.02 & 0.95 \end{bmatrix}$$

  This is a symmetric matrix. The first row represents the interaction of white matter with each of the tissue classes. The second row represents the interaction of gray matter with each of the classes. The third row represents the interaction of CSF with each of the classes. The last row represents the interaction of scalp with each of the classes. This matrix was derived from a manually segmented scan by using the criterion repeated in Equation 4.2, and then forcing symmetry in the matrix by copying the upper triangular half of the matrix to the lower triangular half. The values were sufficiently similar in the two halves that there was no noticeable difference in segmentation when the original vs. forced symmetric matrix were used.

- The Gaussian measurement model parameters (the μ and the σ in Equation 4.6) that were used are given in the following table:

Class	Measurement Model Parameter (μ)	Measurement Model Parameter (σ)
White Matter	101	2
Gray Matter	74	2
CSF	4	3
Scalp	212	4

These measurement models were constructed by fitting individual Gaussian distributions to the intensities of voxels in each class from manually segmented training data. These numbers reflect what is well known to those familiar with brain scans and the gradient echo imaging sequence: the scalp generates the strongest (brightest) signal with a mean intensity of 212 in an image with a normalized range from 0 to 255, and has the highest variance because the class we refer to as "scalp" is comprised of very bright subcutaneous fat, as well as less bright skin. It is followed in brightness by white matter which has a mean intensity 101, gray matter with a mean intensity 74, and then CSF which generates essentially no signal and hence has its mean intensity as 4.

- Stationary priors p^{stat} of Equation 4.4 were measured from manually segmented training data using Equation 4.5. The values used in the experiment are given in the table below:

Class	Stationary Prior
White Matter	0.19
Gray Matter	0.21
CSF	0.5
Scalp	0.1

These priors indicate that white matter and gray matter occur in approximately equal volume in the scan (hence the similar priors), the class CSF, which includes fluid, bone, and air, covers about half of the image (hence the 0.5 prior), and the class scalp occupies the least of the volume (0.1).

- The initial value of $\overline{w}$ was computed by setting J to the zero matrix in Equation 4.1). This is equivalent to the E-Step of the EM-Segmentation algorithm. Subsequent iterations used the value of J specified above.
- We fixed the number of iterations of EM to be 10 in these experiments. This number was empirically derived. In our experience, convergence, in the sense of less than 1% change in the classification between iterations, is typically achieved in 6-7 iterations.
- For each iteration of EM, we performed a single iteration of the Mean-Field solution in the moderately contaminated case, and two iterations in the severely contaminated case.
- The computation time was approximately 4 seconds per slice per EM iteration in a Java implementation of the algorithm that was executed on a single processor of a high performance Sun workstation with 5GB RAM.

4.1.5 Results

We describe the results of this experiment in three steps:

- First we show the results of EM-MF segmentation on a severely contaminated image.
- Then we show results of EM-MF segmentation on a moderately contaminated image and a severely contaminated image.
- Then we compare the results of EM-MF segmentation and EM-Segmentation with manual segmentation for ten MRI scans.

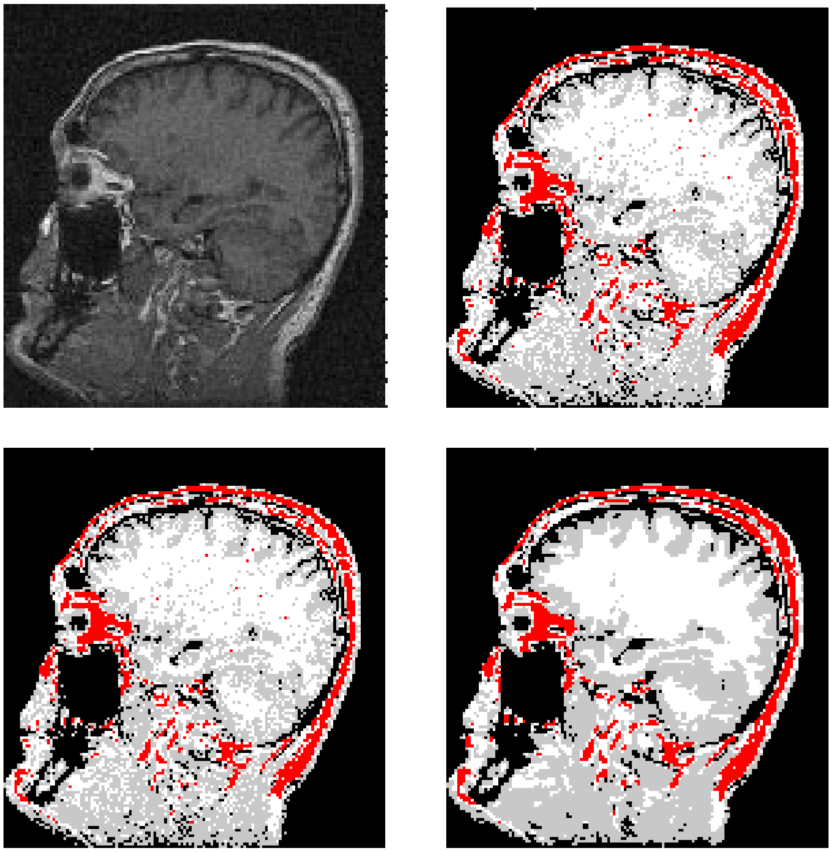

Figure 4-1: An experiment illustrating the robustness of EM-MF segmentation to severe thermal noise in MRI. The top left image is the contaminated input image, and the bottom right image is its EM-MF segmentation. The top right image is a MAP classification of the grayscale image, and is provided as a baseline for visual comparison for the reader. The bottom left image is the EM segmentation of the image. Notice that the EM-MF segmentation is much less fragmented compared to the other segmentations, while preserving the details in the boundaries.

EM-MF Segmentation with Severe Contamination

Figures 4-1, 4-2, and 4-3 show this result. Figure 4-1 shows one slice of the severely contaminated MRI scan in the top left corner. The top right corner of the figure shows a MAP classification of the input image. This is provided as a baseline for visual comparison for those who are familiar with this standard statistical classification technique. The bottom left image shows the EM-Segmentation of the input image. Notice that the thermal noise in the input image leads to misclassifications in the white matter region. The fragmented specks of gray matter in the middle of the white matter volume are misclassifications of white matter. The bottom right image shows the EM-MF segmentation. This image is much less fragmented compared to the EM segmentation, and compared to the MAP segmentation, as is seen in the smooth white matter segmentation. Notice that the details in the segmentation of the white-gray boundaries have been preserved i.e. the fragmentation has not been reduced at the expense of sharpness of the boundaries.

Additional insight may be gained on the impact of the Mean-Field prior on segmentation by contrasting the posterior probabilities that are generated when it is used (EM-MF) vs. when an independent prior is used (EM). We illustrate the differences between these probabilities in Figures 4-2 and 4-3.

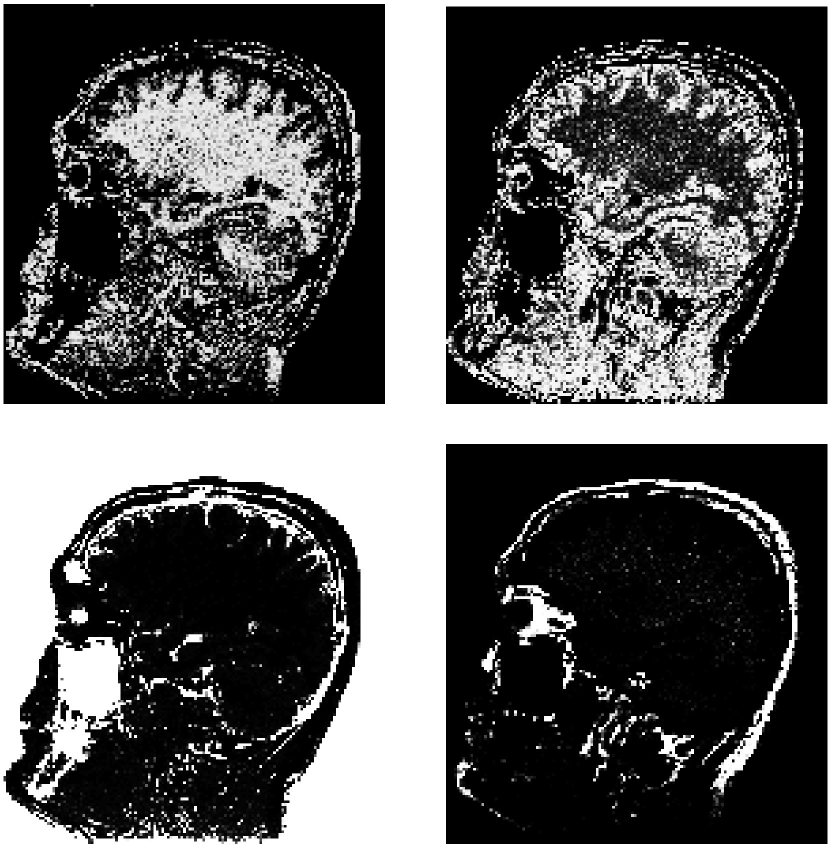

Figure 4-2: These images show the posterior probabilities that result upon convergence of the EM segmentation algorithm when a stationary, independent prior is used. The top left image shows the posterior probabilities for white matter, the top right shows the probabilities for gray matter, the bottom left for CSF, and the bottom right for scalp. These probabilities are thresholded to obtain the EM segmentation shown in the previous figure. The purpose of these images is to contrast them with the posterior probabilities generated by the EM-MF segmentation algorithm in the next figure.

Figure 4-2 shows the posterior probabilities that result upon convergence of the EM-segmentation algorithm. In each of the four images in this figure, each voxel is a real number between 0 (black) and 1 (white), and denotes the posterior probability of a particular tissue class at that voxel. The top left image shows the posterior probabilities for white matter, top right for gray matter, bottom left for CSF (recall that "CSF" refers to the air, the fluid and the bone in the image, which is why, for example, the entire background/air in the image is assigned probability 1 for this class), and bottom right for scalp. These probabilities are the "soft" segmentation that result when EM segmentation converges. The segmentation shown in the bottom left of Figure 4-1 is created by using a winner-take-all (MAP) criterion on these probabilities, i.e. each voxel gets assigned the label of the class with the highest probability from the four images. Notice in the figure that these probabilities are quite fragmented, hence it is no surprise that the resulting segmentation in the previous image is fragmented.

Next we show the unfragmented posterior probabilities when the EM-MF segmentation converges. The images are in Figure 4-3.

As with the previous case, the top left image shows the posterior probabilities for white matter, top right for gray matter, bottom left for CSF, and bottom right for scalp. The EM-MF segmentation shown in the bottom right image of Figure 4-1 is created by using a winner-take-all criterion on these probabilities. Notice that these probabilities are much less fragmented compared to those for EM segmentation in the previous figure, which is why the resulting segmentation in Figure 4-1 has no fragmentation.

EM-MF Segmentation with varying Noise Levels

We use Figures 4-4 and 4-5 to show the performance of EM-MF segmentation of an image that has been contaminated with varying levels of additive noise.

In Figure 4-4, the top left image is a slice from an MR scan. The top right image is the manual segmentation of white matter and gray matter in the image. Note that the manual segmentation has already been "descalped" i.e. all non-brain tissue has been

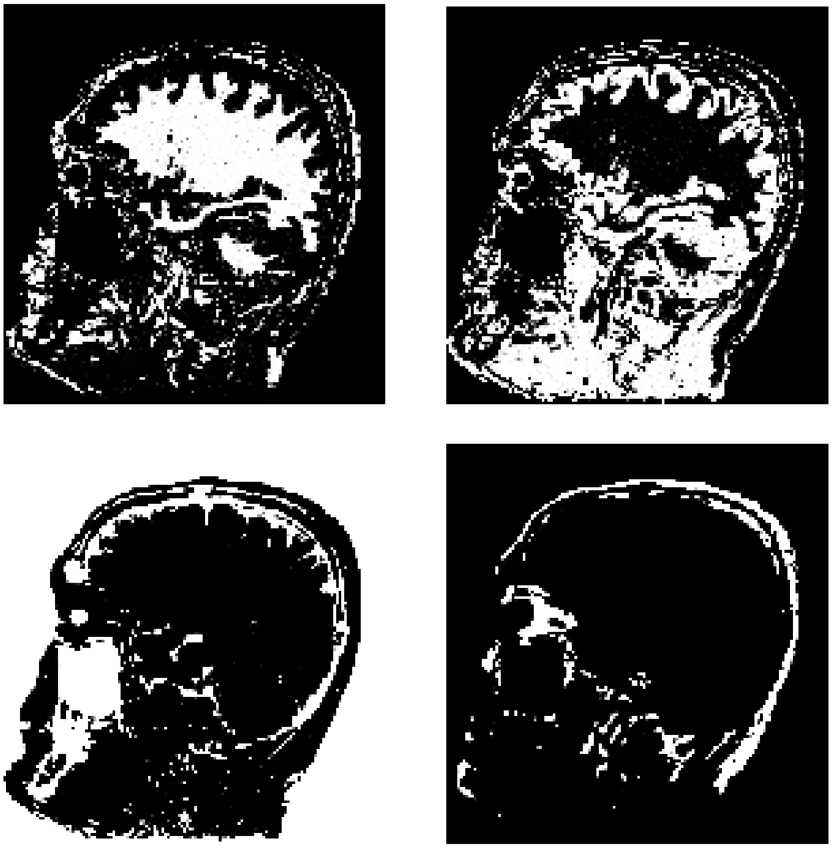

Figure 4-3: These images show the posterior probabilities that result from the EM-MF segmentation algorithm. The top left image shows the posterior probabilities for white matter, the top right shows the probabilities for gray matter, the bottom left for CSF, and the bottom right for scalp. These probabilities are thresholded to obtain the EM-MF segmentation shown in Figure 4-1. Notice that these probabilities are much less fragmented compared to those for EM segmentation in the previous figure, which leads to the reduced fragmentation in the segmentation that we saw in Figure 4-1.

removed from consideration. The bottom left image is a contamination of the MR scan with moderate noise ($\sigma = 5\%$) and the bottom right image is a contamination with severe noise ($\sigma = 10\%$).

Figure 4-5 illustrates the results of EM-Segmentation and EM-MF segmentation on the contaminated images in the bottom row of Figure 4-4. In the top row, the left image is the EM-segmentation of the moderately contaminated image, and the right image is its EM-MF segmentation. In the bottom row, the left image is the EM-segmentation of the severely contaminated image, and the right image is its EM-MF segmentation. Notice that both images in the right column, the EM-MF segmenta-

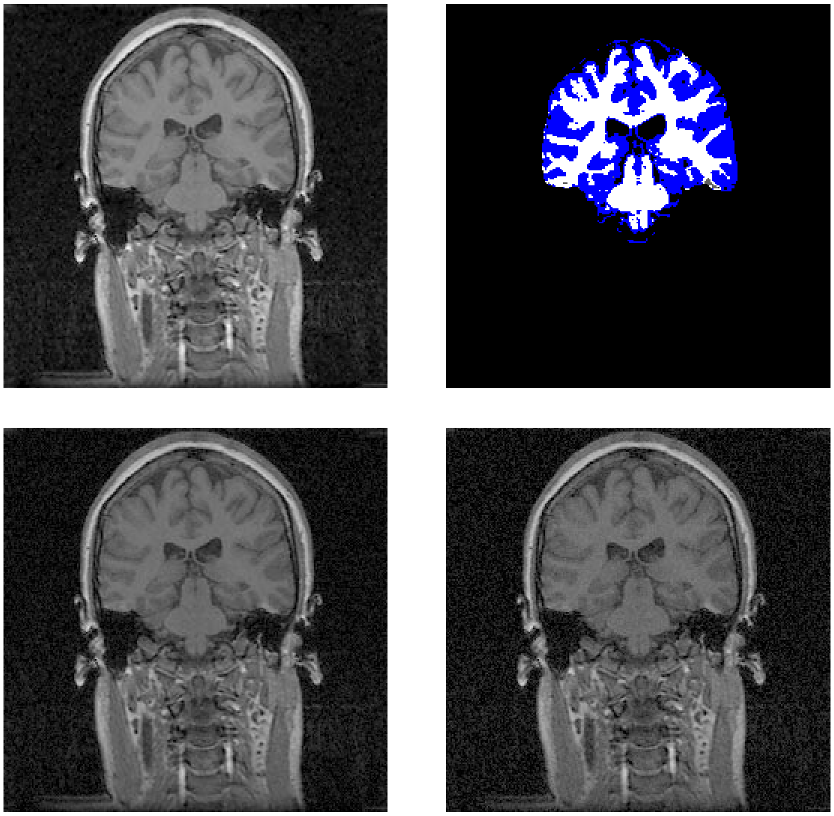

Figure 4-4: The top left image is a slice from an MR scan. The top right image is the manual segmentation of white matter and gray matter in the image. The bottom left image show the moderate contamination of the input, and the bottom right shows the severe contamination.

tions are significantly less fragmented than the corresponding EM-segmentations in the left column. Also notice that the level of fragmentation in the EM-segmentation increases as the contamination of the input image increases, while EM-MF is more robust to the added noise.

Comparison of EM-MF Segmentation and EM-Segmentation with Manual Segmentation

We repeated the previous experiment, of running EM-segmentation and EM-MF segmentation on both moderately and severely contaminated scans, for ten MRI scans. For each scan, we selected five slices, and compared the results of EM-segmentation and EM-MF segmentation against manual segmentation. The manual segmentation of the scans was performed in the original, non-contaminated scans. The comparison

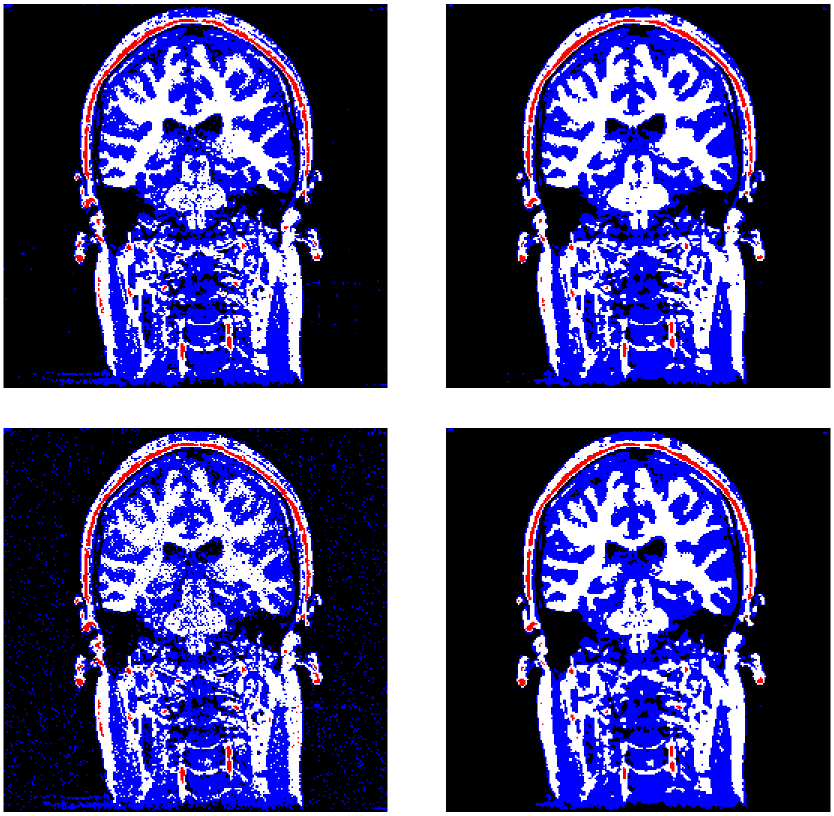

Figure 4-5: In the top row, the left image is the EM-segmentation of the moderately contaminated image, and the right image is its EM-MF segmentation. In the bottom row, the left image is the EM-segmentation of the severely contaminated image, and the right image is its EM-MF segmentation. Notice that both the images in the right column, the EM-MF segmentations, are significantly less fragmented than the corresponding EM-segmentations in the left column.

consists of computing the percent disagreement in classification of the segmentations compared to manual segmentation. It is summarized in the following table:

# #	% Error EM ($\sigma = 10$)	% Error EM-MF ($\sigma = 10$)	% Improv.	% Error EM ($\sigma = 20$)	% Error EM-MF ($\sigma = 20$)	% Improv.
1	10	8	2	14	9	5
2	14	9	5	20	13	7
3	13	10	3	18	11	7
4	12	8	4	17	10	7
5	13	9	4	19	11	8
6	13	10	3	18	11	7
7	10	9	1	19	13	6
8	13	10	3	18	12	6
9	12	8	4	17	10	7
10	14	9	5	21	13	8

The first two columns show the percentage of voxels in EM-Segmentation and EM-MF Segmentation respectively, of moderately contaminated images, that disagree in classification with manual segmentation. The third column is the difference between the numbers in the first two columns, and shows the percent improvement in error in EM-MF compared to EM. Notice that EM-MF segmentation reduces the error by $2-5\%$. The next two columns show the percentage of voxels in EM-Segmentation and EM-MF Segmentation, respectively, of severely contaminated images, that disagree in classification with manual segmentation. Notice that when the noise is increased, the improvement offered by EM-MF over EM is also increased, and is $5-8\%$.

4.1.6 Summary of Experiment and Results

In this experiment we demonstrated the performance of the EM-MF segmentation method, and showed that the introduction of the Mean-Field prior reduces fragmentation in the results while preserving the details in the boundaries. This improves both the visual appeal of the results and the accuracy of the segmentation, as was seen in the above images and the percent errors compared to manual segmentation.

4.2 Geometric Models with EM-MF Segmentation for Brain MRI

In this experiment, we construct a geometric distance model to encode the spatial relationships between structures in the human head. We show how this distance model can be used in conjunction with EM-MF to improve the disambiguation between the scalp and brain tissue.

4.2.1 Motivating Application

We mentioned in Chapter 1 that brain segmentation is typically formulated as a pipelined process in which intensity-based classification is followed by semi-automatic post-processing. The purpose of this post-processing is to "descalp", i.e., remove from consideration as brain tissue, the scalp, and other muscles and fibers that are classified as white or gray matter due to the overlap in the intensities of the materials. Note in the results of EM-MF segmentation in the previous experiment that the scalp is erroneously classified as brain tissue.

The motivation for this experiment is to use a geometric model to capture the information that is necessary for descalping, and the integrate the geometric model as a spatially varying prior into EM-MF segmentation. The result being a segmentation algorithm that simultaneously exploits knowledge of intensity and relative geometry of structures.

4.2.2 The Input Data

The input is gradient Echo MRI scans with $256 \times 256 \times 128$ voxels, and the voxel size is $1mm \times 1mm \times 0.95mm$ (width×height×depth).

The algorithm is trained on three scans that are labeled for white matter, gray matter, CSF, ventricles, and skin. Ventricles are a deep-seated fluid filled structure in the brain, and the skin is the outer boundary of the scalp.

4.2.3 The Experiment

The two structures that we would like to distinguish between, using geometric information, are the scalp and brain tissue. We observe that

- The skin and the ventricles are easy to segment in our images.
- The scalp lies close to the skin and far from the ventricles, while brain tissue lies closer to the ventricles and further from the skin.

Hence we select the skin and the ventricles as primary structures, the scalp, brain tissue, and CSF as secondary structures, and construct a geometric distance model that expresses the spatial relationship between these primary and secondary structures. The geometric model is a two-dimensional distribution of distances to the skin and distances to the ventricles.

Given a new image, we segment primary structures in it, use the distance model to create spatially varying priors on the secondary structures, and use those instead of the spatially stationary priors in EM-MF segmentation.

4.2.4 Details of the Implementation

We implemented the above-mentioned strategy as follows:

- We used the training data to construct the distance models for this experiment in three steps:
    1. For each secondary structure, we computed the distances ρ_{s,l'_1} for all voxels s in the secondary structure to the first primary structure, the skin, using the Chamfer distance transform [Borgefors1986]. A distance transform can be thought of as an image in which the intensity value at each voxel is its distance – approximate Euclidean in our case – from the closest "on" voxel in a given labeled image. The Chamfer algorithm is an efficient method

for computing a distance transform in two passes over an image. We also computed the distances ρ_{s,l'_2} for all voxels s in the secondary structure to the second primary structure, the ventricles in the same way.

2. We jointly histogrammed the distance pair samples computed above for each secondary structure to estimate the respective two-dimensional class conditional density functions

$$P(\rho_{s,l'_1},\rho_{s,l'_2}|s \in BrainTissue), \quad P(\rho_{s,l'_1},\rho_{s,l'_2}|s \in Scalp), \quad P(\rho_{s,l'_1},\rho_{s,l'_2}|s \in CSF).$$

3. We used Bayes' rule to construct the distance model for each secondary structure from the class conditional densities:

$$P(s \in \mathrm{BrainTissue}|\rho_{\mathrm{s,l'_1}},\rho_{\mathrm{s,l'_2}}) = \frac{\mathrm{P}(\rho_{\mathrm{s,l'_1}},\rho_{\mathrm{s,l'_2}}|\mathrm{s} \in \mathrm{BrainTissue})\mathrm{P}(\mathrm{s} \in \mathrm{BrainTissue})}{\mathrm{Z}} \tag{4.7}$$

$$P(s \in \mathrm{Scalp}|\rho_{\mathrm{s,l'_1}},\rho_{\mathrm{s,l'_2}}) = \frac{\mathrm{P}(\rho_{\mathrm{s,l'_1}},\rho_{\mathrm{s,l'_2}}|\mathrm{s} \in \mathrm{Scalp})\mathrm{P}(\mathrm{s} \in \mathrm{Scalp})}{\mathrm{Z}} \tag{4.8}$$

$$P(s \in \mathrm{CSF}|\rho_{\mathrm{s,l'_1}},\rho_{\mathrm{s,l'_2}}) = \frac{\mathrm{P}(\rho_{\mathrm{s,l'_1}},\rho_{\mathrm{s,l'_2}}|\mathrm{s} \in \mathrm{CSF})\mathrm{P}(\mathrm{s} \in \mathrm{CSF})}{\mathrm{Z}} \tag{4.9}$$

where Z is a normalizing constant over all secondary structures given by:

$$\begin{aligned} Z &= P(\rho_{s,l'_1},\rho_{s,l'_2}|s \in \mathrm{CSF})\mathrm{P}(\mathrm{s} \in \mathrm{CSF}) \\ &+ P(\rho_{s,l'_1},\rho_{s,l'_2}|s \in \mathrm{Scalp})\mathrm{P}(\mathrm{s} \in \mathrm{Scalp}) \\ &+ P(\rho_{s,l'_1},\rho_{s,l'_2}|s \in \mathrm{BrainTissue})\mathrm{P}(\mathrm{s} \in \mathrm{BrainTissue}) \end{aligned}$$

We combined the different training images by normalizing the distances from each training image such that their minimum and maximum values were the same.

- Given a novel image, we segmented the primary structures using the following low-level image processing. We segmented the skin by thresholding of the original grayscale image, followed by a connected component analysis to find the largest connected component in the image. The outer boundary of this connected component was marked as the skin. This set of operations is fairly standard practice in the community for segmenting the skin boundary from MR images. The ventricles are fluid filled (dark) structures, and were segmented by region growing from seed points provided by a user.

- For each voxel in the novel image, we computed its distance to the skin, and to the ventricles. Both distances were computed using the Chamfer distance algorithm. We used these distances to look-up the geometry based probabilities of that voxel in the distance model. We replaced the spatially stationary prior in EM-MF segmentation with these spatially varying probabilities.

4.2.5 Results

Figures 4-6, 4-7, and 4-8 show the results of this experiment. Figure 4-6 shows one slice of the input grayscale scan in the top left corner. In the top right corner of the same figure, we show a familiar MAP classification of this input image as a baseline for visual comparison for the other segmentations.In the bottom left image in the figure, we show a pure intensity based EM-Segmentation of the input image. Notice that the scalp is misclassified as white/gray matter in this image. This is simply because there is overlap in the intensity characteristic of these structures, and it is not possible for an algorithm that has uses only intensity models to distinguish them. The bottom right image shows the EM-MF segmentation in conjunction with the above-mentioned geometric distance model. Notice in this image that the scalp areas are no longer confused with white/gray matter. (There is also reduced fragmentation in the white matter in the brain stem area due to the Gibbs prior.)

We show the posterior probabilities for the EM-Segmentation algorithm, and the EM-MF segmentation algorithm with the geometric model in Figures 4-7 and 4-8.

Figure 4-7 shows the posterior probabilities that result upon convergence of the EM-Segmentation algorithm which does not use any geometric information. The top left image shows the posterior probabilities for white matter, top right for gray matter, bottom left for CSF, and bottom right for scalp. These probabilities are the "soft" segmentation from which the (hard) segmentation shown in the bottom left of Figure 4-6 is created by using a winner-take-all criterion. Notice in the images that the regions that belong to the scalp are incorrectly assigned high (white) probabilities of being white/gray matter.

In Figure 4-8 we show the posterior probabilities when a geometric model is used in conjunction with the EM-MF algorithm. As with the previous case, the top left image shows the posterior probabilities for white matter, top right for gray matter, bottom left for CSF, and bottom right for scalp. The segmentation shown in the bottom right image of Figure 4-6 is created by using a winner-take-all criterion on these probabilities. Notice that the regions that belong to scalp are now correctly assigned zero (black) probability in the images for white and gray matter.

4.2.6 Summary of Experiment and Result

In this experiment we demonstrated the use of a geometric distance model to encodes the geometric relationship between brain tissue and the scalp, and primary structures such as the skin and the ventricles. We used this geometric model in conjunction with the EM-MF segmentation method, and applied it to the "descalping" problem in segmentation of brain MRI.

In examining the results of this experiment, we can observe that some misclassifications remain. The isolated pieces of bright tissue that occur near the bottom of the brain stem are incorrectly classified as white matter. These belong to a different class, that we have not yet incorporated into our experiments. In order to correctly handles this case, a "muscle" class should be used in future work.

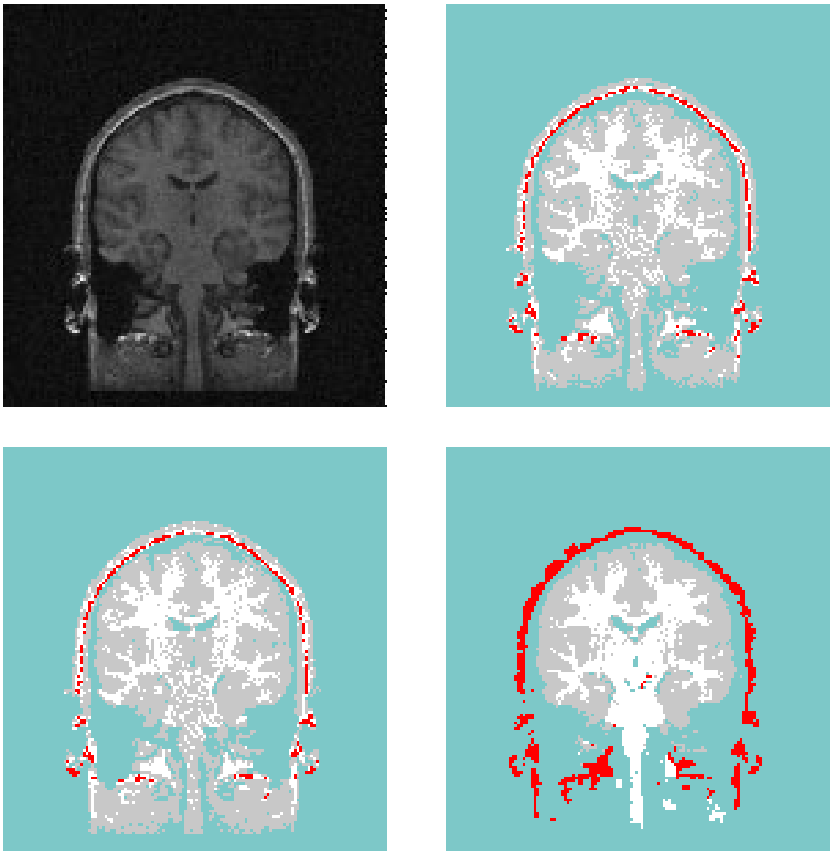

Figure 4-6: An experiment illustrating the use of a geometric distance model in conjunction with the EM-MF algorithm The top left image is the input image. The top right image is a MAP classification of the grayscale image, and is provided as a baseline for visual comparison for the reader. The bottom left image is the EM-MF segmentation without any geometric information, and the bottom right image is its segmentation using a distance model with EM-MF. Notice that using the geometric model removes misclassification of scalp as white/gray matter.

4.3 Geometric Models for Knee MRI

In this experiment, we construct a geometric distance-normal model to encode the spatial relationships between structures in the human knee. We use this distance-normal model in conjunction with a MAP classifier for the segmentation of articular cartilage. Cartilage is a thin structure that shows poor contrast, which makes it impossible to segment using intensity-based classification.

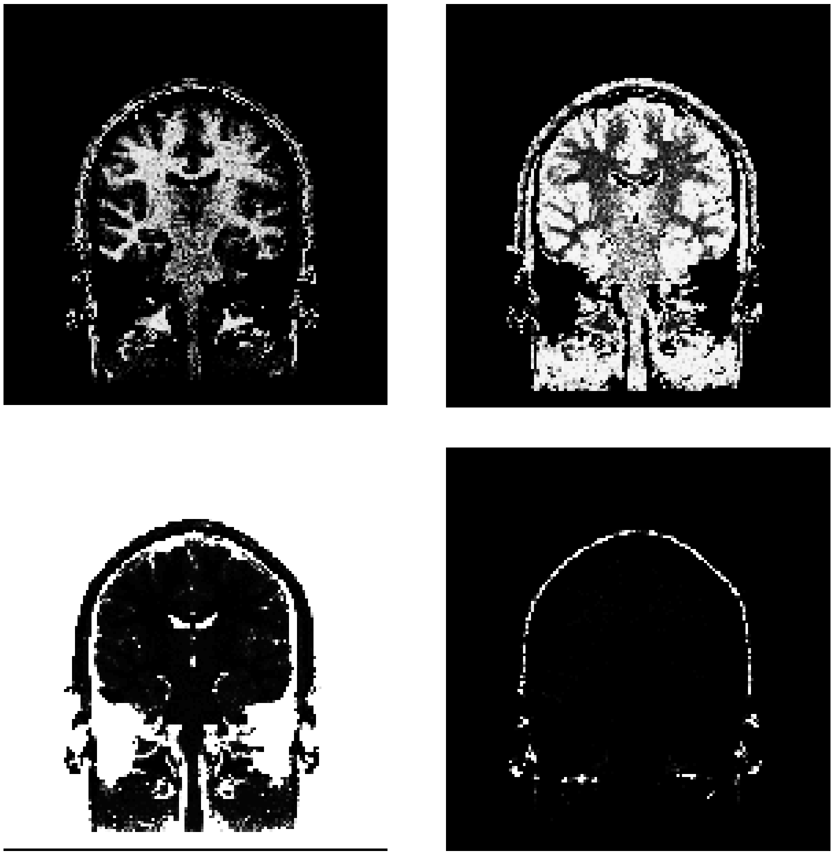

Figure 4-7: These images show the posterior probabilities that result upon convergence of the EM-MF segmentation algorithm when no geometric information is used. The top left image shows the posterior probabilities for white matter, the top right shows the probabilities for gray matter, the bottom left for CSF, and the bottom right for scalp. These probabilities are thresholded to obtain the EM-MF segmentation shown in the previous figure. Notice that the probabilities of white/gray are (incorrectly) high (white) in the regions of the image that belong to scalp.

4.3.1 Motivating Application

The goal of this experiment was to segment a knee MRI into four structures: the femur (thigh bone), the tibia (shin bone), the femoral cartilage (articular cartilage lining the femur) and the tibial cartilage (cartilage lining the tibia). The motivation for segmentation of these structures is to provide patient-specific anatomical models for the arthroscopic surgical simulator being built by Gibson et al. [Gibson et al.1998]. In this simulation system, computer-based knee models are combined with physics-based anatomical modeling, high quality visual rendering, and force-reflective interface devices to provide the surgeon with a virtual testbed for planning and rehearsing

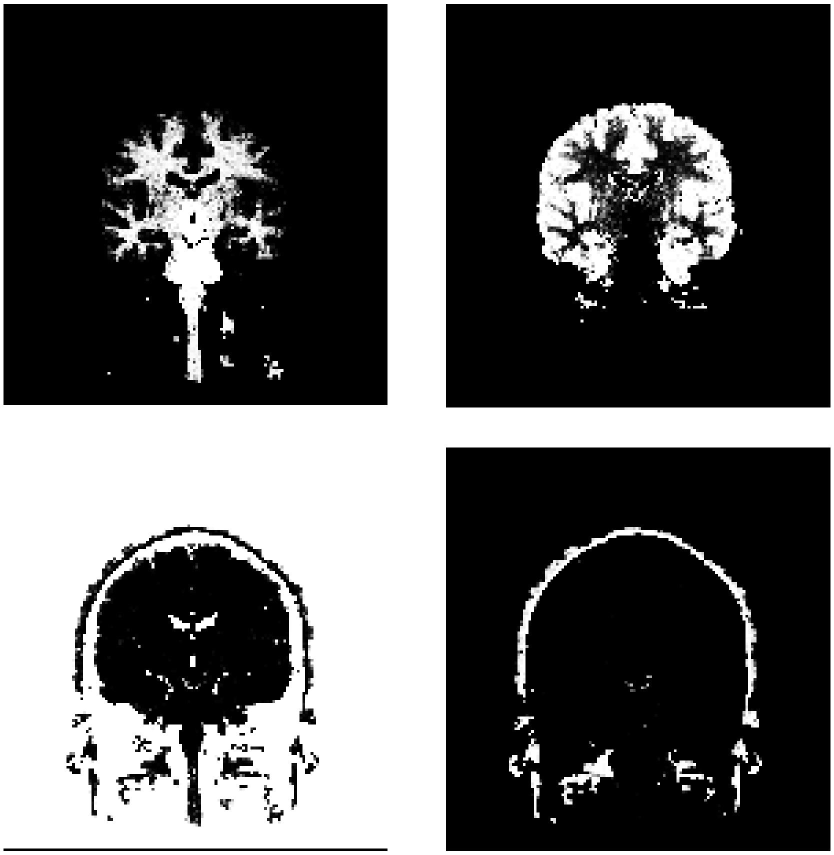

Figure 4-8: These images show the posterior probabilities that result from using a distance model in conjunction with EM-MF. The top left image shows the posterior probabilities for white matter, the top right shows the probabilities for gray matter, the bottom left for CSF, and the bottom right for scalp. These probabilities are thresholded to obtain the bottom right segmentation shown in Figure 4-6. Notice that the probabilities of white/gray are now correctly computed as being very small (black) in the regions of the image that belong to scalp.

patient specific procedures. Current knee models for the surgical simulator have been painstakingly hand-segmented by medical professionals. The hand segmentation has required up to 60 hours per knee model – an effort that would not be economical or practical for patient-specific models.

4.3.2 The Input Data

The input images to be segmented are from a sagittal T2-weighted knee MRI scan with 256×256×124 voxels, and the voxel size is $1mm \times 1mm \times 1mm$ (width×height×depth).

The algorithm is trained on a two-dimensional slices of a labeled scan that is labeled for the four structures: the femur, the tibia, the femoral cartilage, and the tibial cartilage.

4.3.3 The Experiment

In this experiment we were interested in using geometric information to localize articular cartilage in the input images. Note that this experiment in 2D, in contrast with the earlier described experiments, which were all in 3D.

We observe that

- The large bones, the femur and the tibia, easy to segment in our images.
- The location of cartilage is well-defined with respect to the geometry of these bones. Figure 4-9 shows the spatial relationship between the cartilage and the bones. Note that in each case, the cartilage lines the bone.

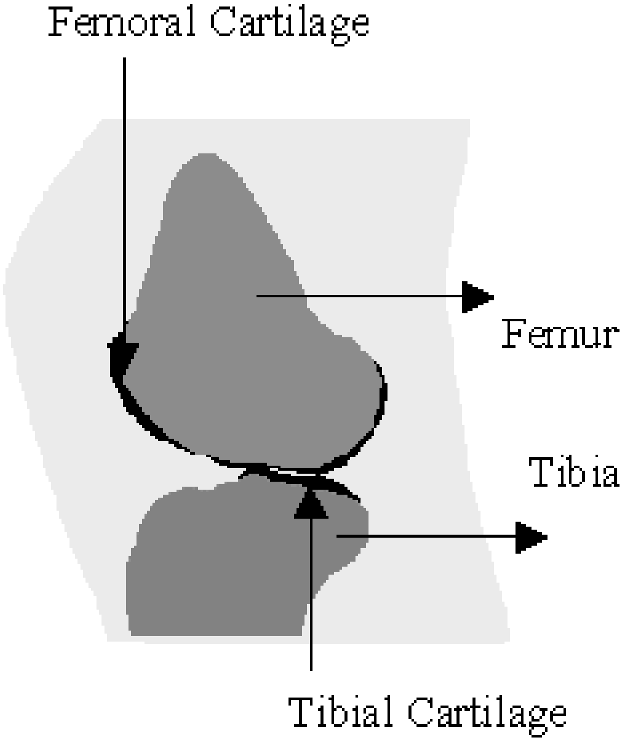

Figure 4-9: A manually segmented slice from a knee MRI which shows four structures: femur, tibia, femoral cartilage and tibial cartilage. Notice how the femoral cartilage lines the lower end of the femur, while the tibial cartilage lines the upper end of the tibia.

Hence we select the femur and the tibia as primary structures, the femoral and articular cartilage secondary structures, and construct geometric distance-normal models that express the spatial relationships between these primary and secondary structures. The geometric model for femoral cartilage is a two-dimensional distribution of distances of the cartilage to the outside bone surface, and the local orientation of the bone where the cartilage connects to it.

Given a new image, we segment primary structures in it, use the distance-normal model to create spatially varying priors on the secondary structures, and use those instead as priors in a standard MAP classifier.

4.3.4 Details of the Implementation

In order to apply the above-mentioned strategy for segmentation of cartilage in knee MRI, we specified: the segmentation algorithm for bones, the method for computing the distances and the normals, the representation used for the class conditional probability distribution functions for cartilage from which the geometric model is derived, and the MAP classification scheme that was used. In this experiment, these details were implemented as follows:

- The segmentation algorithm for bone was motivated by the observation that its appearance in our MRI scans consists of an inner, slightly textured *trabecular* core that is surrounded by a thin, dark layer of *cortical* bone (see Figure 4-10 for a sample slice from this experiment). Using this observation, we segmented each of the bones using an adaptive region-growing algorithm, which took as user input a seed point in the trabecular bone area. We computed the mean and the variance of a local patch around the seed point, and referred to those as the "region mean" and "region variance". We then grew the region, spiraling outward, to include neighboring voxels for which the mean intensity (over a local patch) was within two standard deviations of the region mean. When a voxel was accepted into the region, the region mean/variance were recomputed. Due to the dark, bounding cortical bone, the region growing stopped at the boundary of the trabecular bone. We regularized the boundary of this region using a variant of snakes [Kass, Witkin and Terzopoulos1988]. Results of the region growing for the femur, and the regularized boundary are shown in Figure 4-11.

Figure 4-10: An example slice from the knee MRI used in this experiment. Notice the texture of the trabecular bone, and the fact that it is bounded on the outside with dark cortical bone.

- We computed distances from bone using the Chamfer distance transform algorithm that was mentioned in the previous experiment. We computed the normals to the bone surface using the direction (angle with x-axis) of the local gradients of the bone boundary.

- We estimated the class conditional densities for femoral and tibial cartilage:

$$P(\rho_{s,l'_1}, n_{s,l'_1}|s \in \text{f.Cartilage}), \quad \mathrm{P}(\rho_{\mathrm{s,l'_1}}, \mathrm{n_{s,l'_1}}|\mathrm{s} \in \text{t.Cartilage}),$$

  as two-dimensional histograms, using the distances and normals computed in the previous step. We used fairly coarse bins for the histograms. The bins for

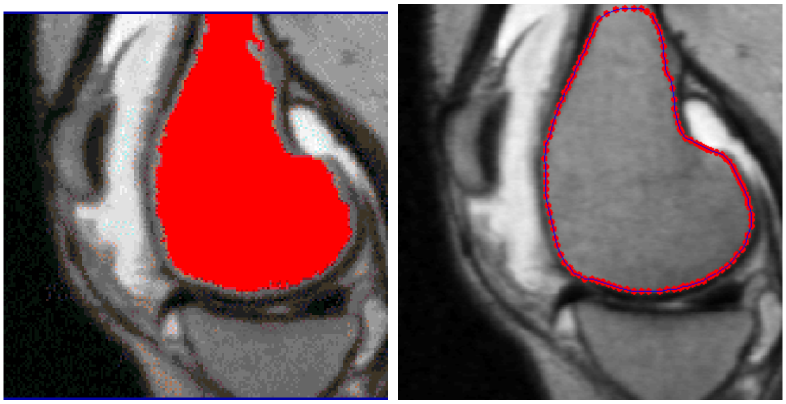

Figure 4-11: Results for segmentation of femur on a knee slice. The left image is the result of region growing on the femur, and and image on the right is the regularized femur boundary.

the distance dimension of the histogram were 5 voxels (mm) wide, and the bins for the normal dimension of the histogram are $\pi/2$ radians wide. This is because all we needed from the normals was whether they were pointing down or up, in order to determine the location of cartilage. The estimated class conditional densities for femoral cartilage and tibial cartilage are shown in Figures 4-12 and 4-13 respectively. Note in the figures that these estimated density functions are localized i.e. have high information content (or low entropy), thereby indicating that the distances and normals are reasonable indicators of the relationship between the bones and the attached cartilage.

- We used the class conditional densities estimated above, in a MAP classifier to segment the femoral cartilage and then tibial cartilage. Given voxel s with intensity I_s, and distance and normals relative to the femur (class label l'_1) ρ_{s,l'_1} and n_{s,l'_1} respectively, we used Bayes' rule to compute the posterior probability of femoral cartilage at s as:

$$\begin{aligned} P(s \in \mathrm{f.cartilage}|\mathrm{\rho_{s,l'_1}, n_{s,l'_1}, I_s}) &= \frac{P(\rho_{s,l'_1}, n_{s,l'_1}, I_s|s \in \mathrm{f.cartilage})\mathrm{P(s \in f.cartilage)}}{Z} \\ &= \frac{P(\rho_{s,l'_1}, n_{s,l'_1}|s \in \mathrm{f.cartilage})\mathrm{P(I_s|s \in f.cartilage)P(s \in f.cartilage)}}{Z} \end{aligned}$$

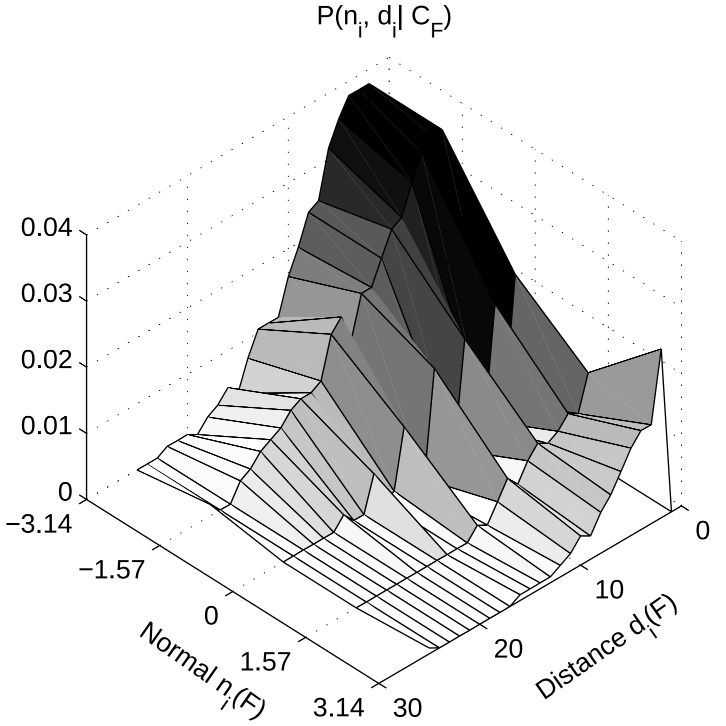

Figure 4-12: Empirical estimate of the joint density for normals and distances for femoral cartilage.

The first term in the numerator on the right hand side is the class conditional density for femoral cartilage that was estimated in the previous step. The second term in the numerator is the measurement model for the intensities of femoral cartilage, which was estimated by fitting a Gaussian distribution to the intensities of femoral cartilage in the training images. The third term in the numerator is the spatially stationary prior probability that a voxel belongs to cartilage. This was computed from a labeled image as a ratio of the cartilage volume to the total volume of the knee in scan. The denominator, Z, is a normalizing term.

A winner-take-all or maximum aposteriori probability (MAP) criterion was used to assign to each voxel the label for femoral cartilage if the posterior probability computed for it according to the above expression was greater than 0.5.

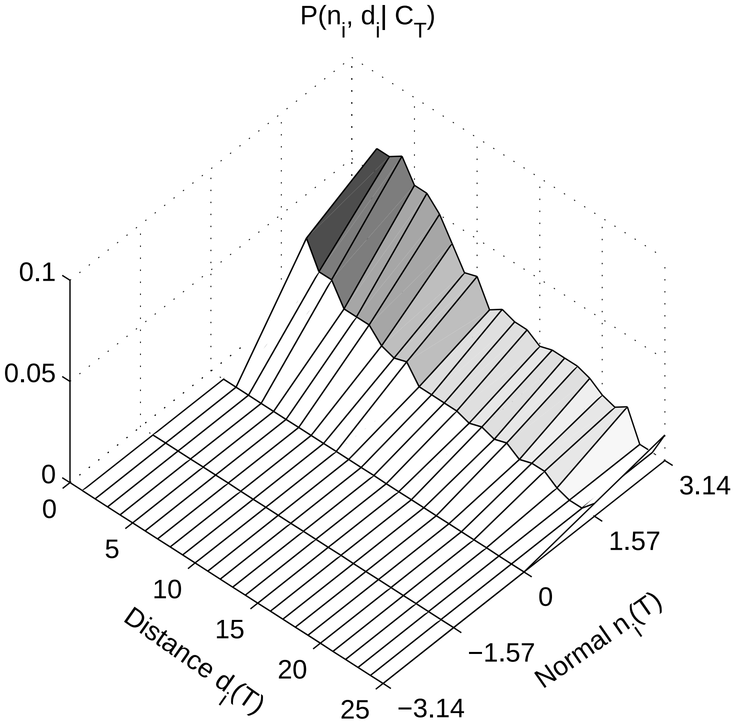

Figure 4-13: Empirical estimate of the joint density for normals and distances for tibial cartilage.

- This experiment was performed on 2D images. We trained and tested our system using “leave k out cross-validation”. In a typical run we trained the system on 60 two-dimensional slices from a single three-dimensional scan and tested on three of the remaining slices. In order to maximize independence between the training and testing images, we left out from the training set the slices that were adjacent to the test images. However, note that all images are from the same 3D scan, hence this independence assumption is not entirely justifiable.

4.3.5 Results

Figures 4-14, 4-15, and 4-16 show the results of this experiment. Figure 4-14 shows three slices of the scan that were left out of the training set in different runs of the experiment. Figure 4-15 shows the segmentation of femoral cartilage generated by our system for these three slices. Figure 4-16 shows the corresponding manual segmentation for visual comparison.

Note that the location of the cartilage generated by our model compares well with manual segmentation. The mean distance between the cartilage segmentation generated by our system and the segmentation generated manually by experts was 1.25 pixels, or 0.32 mm. This distance was computed over three trials of the experiment.

False positives in the first two slices in Figure 4-15 indicate the need for finer quantization of orientation in the estimated joint density of distances and normals. The optimal quantization will require more training data to estimate the joint density (currently we estimate the joint density over 25x5=125 bins using approximately 5000-10000 cartilage points). In order to construct this density, one would train on a set of multiple three-dimensional, registered knee scans.

The segmentation shown above was performed using two classes: femoral cartilage, and rest-of-the-image. In order to use this strategy of segmentation for the complete knee, competing intensity and geometric models for other structures such as the muscle, the ligament, and the meniscus should be used. This would improve the classification, and prevent, for example, some of the false positives in the second image in Figure 4-15.

We used the implementation described above for segmentation of both femoral cartilage and tibial cartilage. Figure 4-17 shows results on tibial cartilage as well as femoral cartilage.

4.3.6 Summary of Knee Experiment

In this experiment we demonstrated the use of geometric distance-normal models to encode the spatial relationship between the bones and the cartilage of the knee. We showed a quantitative evaluation of our results against manual segmentation to show that there is good agreement between our segmentation and manual segmentation. However it should noted that this experiment was performed in 2D while the underlying structure is three-dimensional. The numbers should be interpreted as indicators that this method is promising. It will clearly benefit from additions such as the use of competing models for the geometry of other structures in the knee such as muscles, the ligaments and the meniscus. The assumption of independence between the different primitives of the spatial models (the distances, and the normals) could also be relaxed to allow correlation between them, and capture the varying thickness of cartilage, as is shown in visualizations in [Warfield et al.1998].

4.4 Summary of Experiments and Results

In this chapter we described three experiments in segmentation of MRI. These experiments demonstrate the performance of the EM-MF segmentation method that was presented in Chapter 2, and the Geometric Models that were presented in Chapter 3. In the first experiment we showed the performance of the EM-MF algorithm in reducing fragmentation in the segmentation that is caused due to presence of thermal noise in the scan. In the second experiment we showed how geometric models and EM-MF can be combined to produce a segmentation algorithm that uses knowledge of both intensity as well relative geometry of structures. We applied this segmentation algorithm to "descalp" brain MRI. In the third experiment we showed how to encode the geometric relationship between bone and cartilage to segment articular cartilage in knee MRI by bootstrapping it from the segmentation of the large bones.

Figure 4-14: Three slices from a knee MRI scan. These slices were used to test our cartilage modeling system, and were not included in the set of images used to train the system.

Figure 4-15: Manually generated cartilage segmentation is overlaid in white on the grayscale slices. This figure is presented for visual comparison with the results generated by our system in Figure 4-16.

Figure 4-16: The cartilage segmentation generated by our system is overlaid in white on the grayscale slices. Note that in all cases the cartilage is well localized. False positives in the middle slice (compare with corresponding manual segmentation in Figure 4-15) are discussed in the text.

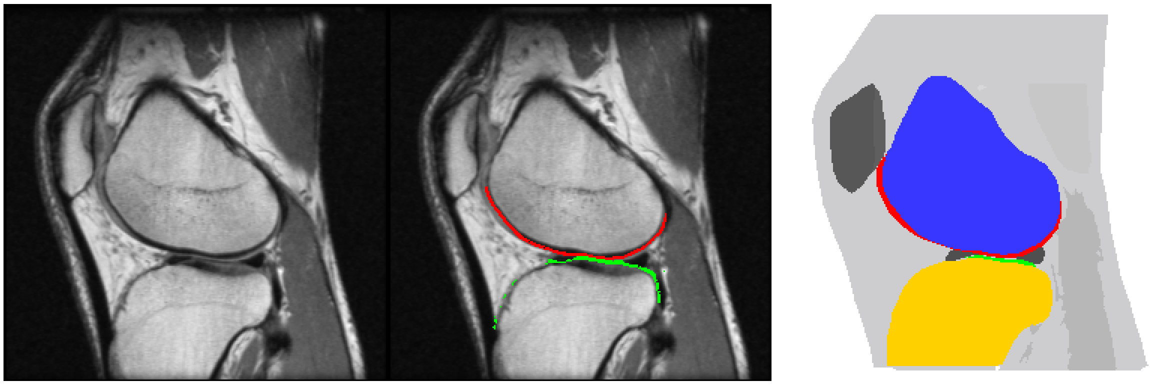

Figure 4-17: Segmentation of femoral cartilage as well as Tibial cartilage generated using our algorithm. The left image is the input grayscale image. The middle image is the segmentation generated by our algorithm – femoral cartilage is shown in green and tibial cartilage is shown in red. The right image is the manual segmentation of cartilage shown for comparison.

Chapter 5

Conclusions and Directions for Future Research

In this chapter we summarize the contributions of this thesis and suggest directions for future research.

5.1 Contributions

This work was driven by the fundamental belief that segmentation of medical imagery is a task that should be attacked with maximal domain knowledge. Specifically, we proposed that four different kinds of knowledge are likely to benefit the segmentation process: *intensity models* that describe the graylevel appearance of the structures of interest, *imaging models* that capture the characteristics of the underlying imaging process, *geometric models* to express the geometric relationships between structures, and *shape models* that describe the shapes of individual structures.

We proposed a statistical framework for segmentation of MRI images, which we call EM-MF. In this EM-MF segmentation framework, we used an imaging model to account for the inhomogeneities due to the MR acquisition process. We used a Gibbs model on tissue class to ensure piecewise homogeneity of the resulting segmentation. We introduced the concept of geometric models for encode spatial relationships between structures, and demonstrated the use of such models in conjunction with

EM-MF segmentation to obtain an algorithm that simultaneously exploits knowledge of the MRI imaging process, intensity models for the appearance of the underlying tissues in the image, as well as the spatial relationships between structures. We demonstrated our techniques in two application areas: segmentation of knee MRI, and segmentation of brain MRI.

5.2 Directions for Future Research

The work presented in this thesis illustrates the use of different kinds of knowledge simultaneously in segmentation of medical imagery. Extensions of the basic ideas presented here are possible both at a theoretical level as well as an application level. At a theoretical level, we suggest the following extensions:

- We presented two examples of geometric models, the *distance model* and the *distance-normal* model. Additional models can be created along the same lines by incorporating more information about the local characteristics of primary structures. For example, a *distance-normal-curvature* model can be created that would measure the principal curvatures of local patches on the surface of the primary structures and encode relationships such as "A is within 10mm of that bump on the lower side of B". [Westin and Kikinis1998] demonstrated the use of tensor models in the segmentation of thin bone from CT. Such models could form the basis for additional primitives to our library of geometric models.

- We presented the notion of *primary* and *secondary* structures in Chapter 3, and gave some information-theoretic heuristics for selecting them manually. Such heuristics could be automated for optimal selection of these structures from labeled training data. For example, a "selection" algorithm could compute the entropy and mutual information of each of the geometric models (distance, distance-normal, distance-normal-curvature etc.) between different structures and select the ones with the most information (least entropy).

- In this work we illustrated *intensity models* (Gaussian, Gibbs), an *imaging model* (slowly varying gain artifact due to inhomogeneity of MRI imaging equipment), and *geometric models* (distance, distance-normal). However, *shape models* remain an important kind of knowledge about anatomical structures ("the brain stem is a tube-like structure") that this work did not venture into. It would be immensely useful to have a shape prior for secondary structures which could be used in conjunction with the relative geometric priors in EM-MF.

At a more application-oriented level, we would suggest the following extensions to facilitate in particular the task of brain segmentation:

- In brain MRI segmentation, we have presented a "descalping" algorithm that uses a geometric distance model in conjunction with EM-MF. This algorithm could be made more robust by using a *coordinate frame model* which would measure the vectors (x, y, z) between the surfaces of the primary and secondary structures, instead of the "lossy" distance model in which the magnitude of the vector from the origin to $(x,\ y,z)$ is stored while the actual coordinates are discarded. It should be kept in mind that (a) the dimensionality of this new model will become three times that of the distance-model, and hence it must be trained on an adequate sized data set, and (b) data sets will need to be rigidly aligned prior to constructing the model as well as prior to segmentation of a novel scan. [Warfield and Kikinis1998] proposes an "Atlas Modulated Segmentation" scheme for segmentation that is of similar flavor.

- In additional preliminary experiments, we have found that using individual geometric models for CSF, ventricles, and background allows us to separate these classes in gradient echo images in which the intensity characteristics of all three are very similar. This leads us to believe that geometric models may be well-suited to the segmentation of deep seated graymatter structures such as the caudate nuclei that are located near the tip of the ventricles, and difficult to identify based on their intensity alone.

In conclusion, a few observations about medical image segmentation:

- It has been our experience that formulating the segmentation task in a Bayesian framework allows us to encode the different kinds of knowledge that are relevant to the problem as prior distributions which can be used simultaneously in an elegant manner. We prefer it to pipelined formulations that violate the principle of least commitment [Marr1982].

- While there is loss of information in the geometric models that we have suggested, we would venture to say that the kinds of geometric relationships that anatomy experts usually use to describe structures are fairly simple: "A is next to B", "A is above B", "A is outside B", "A is inside B", and adequately captured by simple geometric models.

- Deformable models are commonly used in segmentation of medical images. Applying those involves solving a large multidimensional optimization problem to arrive at the optimal segmentation. This approach has rarely been successful simply because there is no easy way to encode global constraints of geometry in these methods, and they are tend to meander in a large search space that could be significantly reduced in size by adding domain knowledge. However, it is not clear how domain knowledge of relative geometry should be added in deformable models.

- Medical data sets are typically large – an MRI scan typically consists of 16M voxels. We believe that it makes perfect sense to replace search in a segmentation algorithm with domain knowledge whenever it is possible. It is attractive to imagine general-purpose segmentation algorithms that will automatically learn relevant knowledge from training data. However, given the size of our data sets, the computational burden associated with such general algorithms is likely to render them impractical clinically. The state of the art in medical image understanding, we believe, will be furthered fastest by end-to-end, application-specific segmentation systems.

Appendix A

Mean-Field Equations

Here we present a derivation of the Mean-Field update equations (Equation 2.36) that were used in Chapter 2. It closely follows material in Chapter 3 of [Elfade1993].

Recall that the energy function that we used in our posterior Gibbs model is given by

$$E(W, g) = -\frac{1}{2} \sum_{s} \sum_{r \in N_s} W_s^T J W_r - \sum_{s} W_s^T g_s \tag{A.1}$$

This can be written in as:

$$E(W, g) = -T \sum_{s} W_s^T h_s \tag{A.2}$$

where h_s is defined as:

$$h_s = \frac{1}{T} \left[g_s + \frac{1}{2} \sum_{r \in N_s} J W_r \right]. \tag{A.3}$$

In the rest of this section, we use W_r^* to denote an approximation to the mean value of the random variable W_r. We use h_s^* to denote an approximation to the mean value of h_s:

$$h_s^* = \frac{1}{T} \left[g_s + \frac{1}{2} \sum_{r \in N_s} J W_r^* \right] \tag{A.4}$$

Recall that the partition function is defined as:

$$Z = \sum_W \exp\left(-\frac{1}{T}E(W,g)\right) \tag{A.5}$$

Based on the above approximation for h_s, we define Z^*, an approximation to the partition function

$$Z^* = \sum_W \exp\left(\sum_s W_s^T h_s^*\right) \tag{A.6}$$

In the above expression for the partition function, the random variables at the neighboring sites have been replaced with their mean values. This is the key step in our Mean-Field approximation. The approximated partition function can be rewritten as:

$$Z^* = \sum_W \prod_s \exp(W_s^T h_s^*) \tag{A.7}$$

Or,

$$Z^* = \sum_{W_1}\sum_{W_2}\cdots\sum_{W_N} \exp(W_1^T h_1^*)\exp(W_2^T h_2^*)\cdots\exp(W_N^T h_N^*) \tag{A.8}$$

Additional manipulations lead to

$$Z^* = \left(\sum_{W_1}\exp(W_1^T h_1^*)\right)\left(\sum_{W_2}\exp(W_2^T h_2^*)\right)\cdots\left(\sum_{W_N}\exp(W_N^T h_N^*)\right) \tag{A.9}$$

Or,

$$Z^* = \prod_s \sum_{W_s} \exp(W_s^T h_s^*) \tag{A.10}$$

Next we use the following identity which is proven in the next section of this appendix:

$$\overline{W_{s_i}} = T\frac{\partial}{\partial g_{s_i}}\log Z \tag{A.11}$$

Based on this identity, we require that the following set of equations hold for our approximations:

$$W_{s_i}^* = T\frac{\partial}{\partial g_{s_i}}\log Z^* \tag{A.12}$$

Let us examine the $\log Z^*$ term:

$$\log Z^* = \sum_s \log \left[\sum_{W_s} \exp(W_s^T h_s^*) \right] \tag{A.13}$$

Given that the state space of W_s is the basis set for $\mathcal{R}^M$, (where M is the number of classes in the segmentation problem), we obtain the following:

$$\log Z^* = \sum_s \log \left[\sum_j \exp(h_{sj}^*) \right] \tag{A.14}$$

Differentiating with respect to g_{si}, which is the i^{th} component of the external field at site s, we obtain:

$$\frac{\partial}{\partial g_{s_i}} \log Z^* = \frac{\partial}{\partial g_{s_i}} \log \left[\sum_j \exp(h_{sj}^*) \right] \tag{A.15}$$

Or,

$$\begin{aligned} \frac{\partial}{\partial g_{s_i}} \log Z^* &= \frac{\frac{\partial}{\partial g_{s_i}} \sum_j \exp(h_{sj}^*)}{\sum_j \exp(h_{sj}^*)} \\ &= \frac{\frac{1}{T} \exp(h_{si}^*)}{\sum_j \exp(h_{sj}^*)} \end{aligned} \tag{A.16}$$

Using the above we can rewrite Equation A.12 as follows:

$$W_{s_i}^* = \frac{\exp(h_{si}^*)}{\sum_j \exp(h_{sj}^*)} \quad . \tag{A.17}$$

This defines a set of consistency equations for our approximation to the mean values of W. The right hand side of the above equation is essentially the "soft-max" of neural network literature [Bishop1995].

It turns out that a solution to this set of equations may be obtained by starting with an initial estimate and iterativly re-evaluating the above equation until convergence is obtained. Convergence properties of such iterations are discussed in [Elfade1993].

In similar applications, the saddle-point approximation can be used to show that estimates obtained by solving such equations are in fact useful approximators of the mean values [Elfade1993, Li1995].

A.1 Proof of Identity $\overline{W}_{si} = T\frac{\partial}{\partial g_{si}} \log\ Z$

We expand the expression on the right hand side as follows:

$$T\frac{\partial}{\partial g_{si}} \log\ Z = \frac{T}{Z}\frac{\partial}{\partial g_{si}} Z \tag{A.18}$$

Rewriting using the definition of the partition function $Z = \sum_W \exp\ (-\frac{1}{T}E(W))$:

$$\begin{aligned} T\frac{\partial}{\partial g_{si}} \log\ Z &= \frac{T}{Z}\frac{\partial}{\partial g_{si}}\sum_W \exp\ (-\frac{1}{T}E(W)) \\ &= \frac{T}{Z}\sum_W [-\frac{\partial}{\partial g_{si}}\frac{E(W)}{T}]\ \exp\ (-\frac{1}{T}E(W)) \end{aligned} \tag{A.19}$$

Using $\frac{\partial}{\partial g_{si}}E(W) = -W_{si}$, we get:

$$\begin{aligned} T\frac{\partial}{\partial g_{si}} \log\ Z &= \frac{1}{Z}\sum_W W_{si}\ \exp\ (-\frac{1}{T}E(W)) \\ &= \sum_W W_{si}\frac{1}{Z}\exp\ (-\frac{1}{T}E(W)) \\ &= \sum_W W_{si}P(W) \end{aligned} \tag{A.20}$$

The expression on the right hand side is by definition the mean or expected value of W_{si} over the Gibbs distribution $P(W)$. We denote this expectation using $\overline{W}_{si}$. Hence we obtain the identity

$$T\frac{\partial}{\partial g_{si}} \log\ Z = \overline{W}_{si} \tag{A.21}$$

Bibliography

[Ayache1995] Ayache, N.1995. Medical computer vision, virtual reality and robotics. *Image and Vision Computing*, 13(4):295–313.

[Aylward et al.1994] Aylward, S., Coggins, J., Cizadlo, T., and Andreasen, N.1994. Spatially Invariant Classification of Tissues in MR Images. In *Proceedings of the Third Conference on Visualization in Biomedical Computing*. SPIE.

[Besag1986] Besag, J.1986. On the statistical analysis of dirty pictures. *Proceedings of the Royal Statistical Society*, B-48(3):259–302.

[Bezdek, Hall and Clarke1993] Bezdek, J., Hall, L., and Clarke, L.1993. Review of MR Image Segmentation Techniques using Pattern Recognition. *Medical Physics*, 20(4):1033 – 1048.

[Bishop1995] Bishop, C.1995. *Neural Networks for Pattern Recognition*. Oxford University Press.

[Black and Chaberie1998] Black, P. and Chaberie, A.1998. Clinical experience with xplan. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*.

[Blake and Zisserman1987] Blake, A. and Zisserman, A.1987. *Visual Reconstruction*. MIT Press.

[Borgefors1986] Borgefors, G.1986. Distance transformations in digital images. *Computer Vision, Graphics, and Image Processing*, 34:344–371.

[Brummer et al.1993] Brummer, M., Mersereau, R., Eisner, R., and Lewine, R.1993. Automatic detection of brain contours in MRI data sets. *IEEE Transactions on Medical Imaging*, 12(2):153–166.

[Bucholtz et al.1997] Bucholtz, R., Yah, D., Trobaugh, J., McDurmont, L., Sturm, C., Baumann, C., Henderson, J., Levy, A., and Kessman, P.1997. The correction of stereotactic inaccuracy caused by brain shift using an interaoperative ultrasound device. In *Proceedings First Joint CVRMED/MRCAS*, Grenoble France.

[Chakraborty, Staib and Duncan1994] Chakraborty, A., Staib, L., and Duncan, J.1994. An integrated approach to boundary finding in medical images. In *Proceedings of the IEEE Workshop on Biomedical Image Analysis*, Seattle, WA. IEEE.

[Cline et al.1990] Cline, H., Lorensen, W., Kikinis, R., and Jolesz, F.1990. Three-Dimensional Segmentation of MR Images of the Head Using Probability and Connectivity. *JCAT*, 14(6):1037–1045.

[Colchester et al.1996] Colchester, A., Zhao, J., Holton-Tainter, K., Henri, C., Maitland, N., Roberts, P., Harris, C., and Evans, R.1996. Development and preliminary evaluation of vislan, a surgical planning and guidance system using intra-operative video imaging. *Medical Image Analysis*, 191:73–90.

[Cootes et al.1994] Cootes, T., Hill, A., Taylor, C., and Haslam, J.1994. Use of active shape models for locating structure in medical images. *Image and Vision Computing*, 12(6):355–365.

[Cover and Thomas1991] Cover, T. and Thomas, J. A.1991. *Elements of Information Theory*. John Wiley and Sons, New York.

[Dempster, Laird and Rubin1977] Dempster, A., Laird, N., and Rubin, D.1977. Maximal likelihood form incomplete data via the em algorithm. *Proceedings of the Royal Statistical Society*, B-39:1–38.

[Duda and Hart1973] Duda, R. and Hart, P.1973. *Pattern Classification and Scene Analysis*. John Wiley and Sons.

[Elfade1993] Elfade, I. M.1993. From fields to networks. In *Ph.D. Thesis, MIT.*

[ETH] *Image Science Group.* ETH Zürich. http://www.vision.ee.ethz.ch.

[Ettinger1997] Ettinger, G. J.1997. Hierarchical registration of medical images. In *Ph.D. Thesis, MIT.*

[Faugeras and Berthod1981] Faugeras, O. and Berthod, M.1981. Improving consistency and reducing ambiguity in stochastic labeling: An optimization approach. *PAMI*, 3(4):412–424.

[GE] *Three Dimensional Medical Reconstruction.* GE Corporate Research & Development. http://www.crd.ge.com/esl/cgsp/projects/medical.

[Geiger and Girosi1991] Geiger, D. and Girosi, F.1991. Parallel and deterministic algorithms from mrfs: Surface reconstruction. *PAMI*, 13(5):401–412.

[Geman and Geman1984] Geman, S. and Geman, D.1984. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *PAMI*, 6(6):721–741.

[Gerig et al.1989] Gerig, G., Kuoni, W., Kikinis, R., and Kübler, O.1989. Medical Imaging and Computer Vision: an Integrated Approach for Diagnosis and Planning. In *Proc. 11'th DAGM Symposium*, pages 425–443. Springer.

[Gering1999] Gering, D.1999. A system for surgical planning and guidance using image fusion and interventional mr. In *Masters Thesis, Massachusetts Institute of Technology.*

[Gibson et al.1998] Gibson, S., Fyock, C., Grimson, W., Kanade, T., Kikinis, R., Lauer, H., McKenzie, N., Mor, A., Nakajima, S., Ohkami, H., Osborne, R., Samosky, J., and Sawada, A.1998. Volumetric object modeling for surgical simulation. In *Medical Image Analysis.*

[Gohagan et al.1987] Gohagan, J., Spitznagel, E., Murphy, W., Vannier, M., et al.1987. Multispectral Analysis of MR Images of the Breast. *Radiology*, (163):703 – 707.

[Grimson et al.1995] Grimson, W., Ettinger, G., White, S., Gleason, P., Lozano-Perez, T., Wells III, W., and Kikinis, R.1995. Evaluating and validating an automated registration system for enhanced reality visualization in surgery. In *CVRMed95*.

[Grimson et al.1997] Grimson, W., Ettinger, G. J., Kapur, T., Leventon, M. E., Wells III, W., and Kikinis, R.1997. Utilizing segmented mri data in image guided surgery. *IJPRAI*, 11(8).

[Grimson et al.1994] Grimson, W., Lozano-Perez, T., Wells III, W., Ettinger, G., and White, S.1994. An automatic registration method for frameless stereotaxy, image, guided surgery and enhanced reality visualization. In *CVPR94*, pages 430–436.

[GUYS] *Computational Image Science Group.* Guy's and St. Thomas' Hospital. http://www-ipg.umds.ac.uk/cisg/indexn.htm.

[Held et al.1998] Held, K., Rota Kopps, E., Krause, B., Wells, W., Kikinis, R., and Muller-Gartner, H.1998. Markov random field segmentation of brain mr images. *IEEE Transactions on Medical Imaging*, 16:878–887.

[Highnam1992] Highnam, R.1992. Model-based enhancement of mammographic images. In *Ph.D. Thesis, Computing Laboratory, Oxford University.*

[Hinton] *Anterior Cruciate Ligament.* Dr. Alan Hinton, M.D., Orthopaedic Surgeon. http://lac.laci.net/pweb/hintonmd/acl.htm.

[Hirayasu et al.1998] Hirayasu, Y., Shenton, M., Salisbury, D. F., Dickey, C., Fischer, I., Mazzoni, P., Kisler, T., Arakaki, H., Kwon, J., Anderson, J., Yurgelun-Todd, D., Tohen, M., and McCarley, R.1998. Lower left temporal lobe mri volumes in patients with first-episode schizophrenia compared with psychotic patients with first-episode affective disorder and normal subjects. *American Journal of Psychiatry*, 155(10):1384–1391.

[Horn1986] Horn, B.1986. Robot vision. In *MIT Press.*

[INRIA] *Epidaure Research Project.* INRIA. http://www-sop.inria.fr/epidaure/Epidaure-eng.html.

[Ising1925] Ising, E.1925. Beitrag zur theorie des ferromagnetismus. *Zeitschr. f. Physik*, 31:253–258.

[JHU] *Computer-Integrated Surgical Systems and Technology.* Johns Hopkins University. http://cisstweb.cs.jhu.edu.

[Jolesz et al.1997] Jolesz, F. et al.1997. Image-guided procedures and the operating room of the future. *Radiology*, 204:601–612.

[Kass, Witkin and Terzopoulos1988] Kass, M., Witkin, A., and Terzopoulos, D.1988. Snakes: Active contour models. *IJCV*, 1(4):321–331.

[Kaus et al.1998] Kaus, M., Warfield, S., Jolesz, F., and Kikinis, R.1998. Adaptive template moderated brain tumor segmentation in mri. In *Bildverarbeitung fur die Medizin*, pages 102–106. Springer Verlag.

[Kirpatrick, Gelatt and Vecchi1983] Kirpatrick, S., Gelatt, Jr., C., and Vecchi, M.1983. Optimization by simulated annealing. *Science*, 220:671–680.

[Kohn et al.1991] Kohn, M., Tanna, N., Herman, G., et al.1991. Analysis of Brain and Cerebrospinal Fluid Volumes with MR Imaging. *Radiology*, (178):115 – 122.

[LEUVEN] *Medical Image Computing.* Leuven. http://www.esat.kuleuven.ac.be/mi2/mic.html.

[Leventon1997] Leventon, M.1997. A registration, tracking, and visualization system for image guided surgery. In *Masters Thesis, Massachusetts Institute of Technology.*

[Li1995] Li, S.1995. *Markov Random Field Modeling in Computer Vision.* Springer-Verlag.

[Lorigo et al.1998] Lorigo, L., Faugeras, O., Grimson, W., and Kikinis, R.1998. Segmentation of bone in clinical knee mri using texture-based geodesic active contours. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, pages 1195–1204.

[Marr1982] Marr, D.1982. *Vision.* Freeman.

[Marroquin1984] Marroquin, J. L.1984. Probabilistic solution of inverse problems. In *Ph.D. Thesis, MIT.*

[Martin1995] Martin, J.1995. Characterization of neuropathological shape deformations. In *Ph.D. Thesis, MIT.*

[MAYO] *Biomedical Image Resource.* Mayo Clinic. http://www.mayo.edu/bir/.

[McInerny and Terzopoulos1996] McInerny, T. and Terzopoulos, D.1996. Deformable models in medical image analysis: a survey. *Medical Image Analysis*, 1(2):91–108.

[MERL] *Medical Applications Research.* Mitsubishi Electric Research Laboratory. http://www.crd.ge.com/esl/cgsp/projects/medical.

[Metropolis et al.1953] Metropolis, N., Rosenbluth, A., Rosenbluth, A. N., Teller, A., and Teller, E.1953. Equation of state calculations by fast computing machines. *J. Chem. Phys.*, 21(6):1087–1092.

[MGH] *http://research.neurosurgery.mgh.harvard.edu.* Massachusetts General Hospital and Harvard Medical School.

[MNI] *Image Guided Neurosurgery Laboratory.* Montreal Neurological Institute and the University of Western Ontario. http://www.irus.rri.on.ca/igns.

[Neal and Hinton1998] Neal, R. M. and Hinton, G. E.1998. A view of the em algorithm that justifies incremental, sparse, and other variants. In Jordan, M. I., editor, *Learning in Graphical Models*, pages 355–368. Dordrecht: Kluwer Academic.

[O'Toole1998] O'Toole, R.1998. Assessing skill and learning in surgeons and medical students using a force feedbafck surgical simulator. In *Medical Image Computing and Computer Assisted Intervention (MICCAI).*

[Parisi1988] Parisi, G.1988. *Statistical Field Theory.* Addison-Wesley.

[Pavlidis1982] Pavlidis, T.1982. Algorithms for graphics and image processing. In *Computer Science Press.*

[Perona, Shiotan and Malik1994] Perona, P., Shiotan, T., and Malik, J.1994. Anisotropic diffusion. *Geometry-Driven Diffusion in Computer Vision*, pages 73–92.

[Peterson and Soderberg1989] Peterson, C. and Soderberg, B.1989. A new method for mapping optimization problems onto neural networks. *International Journal of Neural Systems*, 1(1).

[Pham and Prince1999] Pham, D. L. and Prince, J. L.1999. An adaptive fuzzy c-means algorithm for image segmentation in the presence of image inhomogeneities. *Pattern Recognition Letters*, 20(1):57–68.

[Rosenfeld, Hummel and Zucker1976] Rosenfeld, A., Hummel, R., and Zucker, S.1976. Scene labeling by relaxation operations. *IEEE Transactions on Systems, Man, and Cybernetics*, 6(6):420–433.

[Samosky1998] Samosky, J.1998. Simulation of arthroscopic probing in the knee: Ex vivo measurement, tissue modeling and validation studies for an elastic, volumetric, patient-specific model of articular cartilage probing for arthroscopic surgery simulation. In *Ph.D. Thesis Proposal, Division of Health Sciences and Technology, MIT.*

[Serra1982] Serra, J.1982. *Image Analysis and Mathematical Morphology.* London Academic.

[SHADYSIDE] *Medical Robotics and Computer Assisted Surgery.* Shadyside Hospital and Carnegie Mellon University. http://www.mrcas.ri.cmu.edu.

[Solloway et al.1997] Solloway, S., Hutchinson, C., Waterton, J., and Taylor, C.1997. The use of active shape models for making thickness measurements of articular cartilage from mr images. *MRM*, 37(11):943–952.

[SPL] *Surgical Planning Lab.* Brigham & Women's Hospital and Harvard Medical School. http://www.splweb.bwh.harvard.edu:8000.

[Staib and Duncan1992] Staib, L. and Duncan, J.1992. Boundary finding with parametrically deformable models. *PAMI*, 14(11):1061–1075.

[Szekely et al.1996] Szekely, G., Kelemen, A., Brechbuhler, C., and Gerig, G.1996. Segmentation of 2d and 3d objects from mri volume data using constrained elastic deformations of flexible fourier contour and surface models. *Medical Image Analysis*, 1(1):19–34.

[Therrien1984] Therrien, J.1984. *Pattern Recognition.*

[UCLA] *Laboratory of Neuro Imaging.* UCLA. http://www.loni.ucla.edu.

[UCSF] *Department of Radiology.* UCSF. http://www.radiology.ucsf.edu.

[Vandermeulen et al.1996] Vandermeulen, D., Descombes, X., Suetens, P., and Marchal, G.1996. Unsupervised regularized classification of multi-spectral MRI. In *Proceedings of the Fifth Conference on Visualization in Biomedical Computing.* SPIE.

[Vannier et al.1985] Vannier, M., Butterfield, R., Jordan, D., Murphy, W., et al.1985. Multi-Spectral Analysis of Magnetic Resonance Images. *Radiology*, (154):221 – 224.

[Warfield and Kikinis1998] Warfield, S. and Kikinis, R.1998. Adaptive template moderated spatially varying statistical classification. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, pages 231–238.

[Warfield et al.1998] Warfield, S., Winalski, C., Jolesz, F., and Kikinis, R.1998. Automatic segmentation of mri of the knee. In *ISMRM Sixth Scientific Meeting and Exhibition*, page 563.

[Wells III et al.1994] Wells III, W., Kikinis, R., Grimson, W., and Jolesz, F.1994. Statistical Intensity Correction and Segmentation of Magnetic Resonance Image

Data. In *Proceedings of the Third Conference on Visualization in Biomedical Computing*. SPIE.

[Wells III et al.1996] Wells III, W., Kikinis, R., Grimson, W., and Jolesz, F.1996. Adaptive segmentation of mri data. *IEEE Transactions on Medical Imaging*, 15(4):429–442.

[Westin and Kikinis1998] Westin, C. and Kikinis, R.1998. Tensor controlled local structure enhancement of ct images for bone segmentation. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, pages 1205–1212.

[Whitaker and Gerig1994] Whitaker, R. and Gerig, G.1994. Vector-valued diffusion. In ter Haar Romeny, B., editor, *Geometry-Driven Diffusion in Computer Vision*, pages 93–133.

[Xu et al.1999] Xu, C., Pham, D. L., Rettmann, M. E., Yu, D. N., and Prince, J. L.1999. Reconstruction of the human cerebral cortex from magnetic resonance images. *IEEE Trans. Med. Imag.* to appear.

[Zeng et al.1998] Zeng, X., Staib, L., Schultz, R., and Duncan, J.1998. Segmentation and measurement of the cortex from 3d mr images. In *Medical Image Computing and Computer Assisted Intervention (MICCAI).*