Medical Image Processing and Analysis

J. Michael Fitzpatrick Milan Sonka

Example of the LaTeX layout
February 2, 2000

Contents

CHAPTER 1
Brain Image Analysis and Atlas Construction

Paul M. Thompson
University of California, Los Angeles
Michael S. Mega
University of California, Los Angeles
Katherine L. Narr
University of California, Los Angeles
Elizabeth R. Sowell
University of California, Los Angeles
Rebecca E. Blanton
University of California, Los Angeles
Arthur W. Toga
University of California, Los Angeles

1.1 Challenges in Brain Image Analysis

The tremendous pace of development in brain imaging technologies has revolutionized our ability to investigate brain structure and function. Techniques are now available to capture features of anatomy and function at both molecular and whole-brain scales, mapping neuronal dynamics and gene expression as well as growth and degenerative processes that span multi-year time-scales. The number of brain imaging investigations is also increasing exponentially (Fox, 1997). A major goal of these studies is to analyze how the dynamically changing brain varies across age, gender, disease, across multiple imaging modalities, and in large human populations. To tackle these questions, many laboratories are using sophisticated algorithms for brain image analysis. Engineering approaches drawn from computer vision, image analysis, computer graphics and artificial intelligence research fields are required to manipulate, analyze and communicate brain data. Novel image analysis algorithms continue to uncover new patterns of altered structure and function in individuals and clinical populations, and mathematical strategies are

being developed to relate these patterns to clinical, demographic and genetic parameters.

1.1.1 Image Analysis and Brain Atlases

In this chapter, we review current challenges in brain image analysis, focusing on the main algorithms, their technical foundations, and their scientific and clinical applications. The approaches include methods for automated registration and segmentation, anatomical parameterization and modeling, tissue classification and shape analysis, and pathology detection in individuals or groups. Algorithms are also described for generating *digital brain atlases*. Atlases are fundamental to brain image analysis, as they offer a powerful framework to synthesize the results of disparate imaging studies (Roland and Zilles, 1994; Kikinis et al., 1996; Toga and Thompson, 1998). Built from one or more representations of the brain, atlases are annotated representations of anatomy in a 3D coordinate system. They serve as standardized templates on which other brain maps can be overlaid, for subsequent comparison and integration. To align new imaging data with an atlas, a variety of registration algorithms may be employed (see Chapter 8). Once registered, brain maps can be pooled across subjects, and combined mathematically and statistically. As such, atlases provide a standardized 3D coordinate system to express observations from different individuals, and a framework for interlaboratory communication.

1.1.2 Adaptable Brain Templates

Imaging algorithms are also significantly improving the flexibility of digital brain templates. *Deformable brain atlases* are adaptable brain templates that can be individualized to reflect the anatomy of new subjects. This allows the automated labeling of structures in new patients' scans (Evans et al., 1991; Gee et al., 1993; Christensen et al., 1993; Sandor and Leahy, 1995; Rizzo et al., 1995; Toga and Thompson, 1997; Haller et al., 1997). High-dimensional image registration, or warping algorithms (Christensen et al., 1993; 1996; Collins et al., 1994a; Thirion, 1995; Rabbitt et al., 1995; Warfield et al., 1995; Davatzikos, 1996; Thompson and Toga, 1996; Bro-Nielsen and Gramkow, 1996; Gee et al., 1998; Grenander and Miller, 1998; see Toga, 1998 for a review) apply local dilations and contractions to a labeled digital atlas, elastically deforming it to fit a new subject's anatomy. These algorithms can transfer 3D maps of functional and vascular territories onto the scan of any subject, as well as information on tissue types, cytoarchitecture, histologic and neurochemical content (Mega et al., 1997). These algorithms are discussed in detail, later in the chapter.

1.1.3 Mapping Structural Differences

As a valuable by-product, 3D warping algorithms also *quantify* local and global shape changes. The complex profiles of dilation and contraction required to warp an atlas onto a new subject's brain provide an index of the anatomical shape differences between that subject's brain and the atlas (Davatzikos et al., 1996; Thompson et al., 1997; Ashburner et al., 1998). Differences in regional shape can be assessed by the displacement required to locally deform one brain volume into another, and can be further examined by applying vector and tensor field operators to the transformation field (Thompson and Toga, 1998; Thirion et al., 1998; Thompson et al., 2000). As a result, deformable atlases not only adapt to individual anatomy, but they offer a powerful strategy to analyze developmental, age-related or pathologic variations.

1.1.4 Probabilistic Atlases

As imaging studies expand into ever-larger patient populations, *population-based* atlases are required to identify statistical trends. These include group patterns of tissue distribution, anatomical asymmetry, and disease-specific features, which are often obscured in an individual due to the complexity of normal anatomy (Thompson et al., 2000; Mega et al., 2000; Narr et al., 2000a,b). *Probabilistic atlases* identify these patterns by storing detailed quantitative information on cross-subject variations in brain structure and function (Mazziotta et al., 1995). Information stored in the atlas can then be used to identify anomalies and label structures in new patients. Based on well-characterized subject groups, these atlases contain thousands of structure models, as well as composite maps, and average templates, and allow visualization of structural variability, asymmetry and group-specific differences (Fig. 1).

1.1.5 Encoding Cortical Variability

Because cortical anatomy is so variable across subjects, it presents unique challenges in brain image analysis. Specialized algorithms are needed to correct for wide variations in gyral patterns, and to identify alterations in cortical anatomy in groups or individuals. Using mathematical equations derived from continuum mechanics and random field theory, probabilistic atlases can be used to detect cortical features not apparent in individual patients' scans. These include subtle changes in brain asymmetry during late brain development and early changes in gray matter distribution in degenerative disease (Thompson et al., 1997, 2000; Sowell et al., 2000). From an algorithmic standpoint, information on anatomic and functional variability can also guide algorithms for knowledge-based image analysis, including automated image labeling (Collins et al., 1994; Pitiot et al., 2000), tissue

classification (Zijdenbos and Dawant, 1994) and functional image analysis (Dinov et al., 2000).

1.1.6 Disease-Specific Atlases

Once built, population-based atlases can be stratified into subpopulations to reflect the unique anatomy and physiology of a specific group. Atlases are being built to represent populations with Alzheimer's Disease (Fig. 1; Thompson et al., 2000a,b; Mega et al., 2000) and schizophrenia (Narr et al., 2000a,b,c). Because of disease-induced changes in morphology, the anatomy of these populations is not well-accommodated by existing atlases. In dementia studies, for example, registration algorithms can align *post mortem* and *in vivo* data, to correlate the structural, metabolic, molecular and histologic hallmarks of the disease (Mega et al., 1997). The resulting atlases link disease-specific features with demographic factors such as age, gender, handedness, as well as specific clinical or genetic parameters (Mazziotta et al., 1995; Toga and Mazziotta, 1996; Mega et al., 1998).

1.1.7 Dynamic (4D) Brain Data

Finally, new brain atlases are being built to incorporate dynamic data. Despite the significant challenges in expanding the atlas concept to the time dimension, *dynamic brain atlases* are beginning to include probabilistic information on growth rates and changes in tissue distribution to assist research into pediatric disorders (Thompson and Toga, 1998; Sowell et al., 1999a,b, 2000a,b). Registration algorithms that deform one anatomy into the shape of another can also be used to measure dynamic patterns of structural change during brain development, tumor growth, or degenerative disease processes (Toga et al., 1996; Thompson et al., 2000).

All these atlases require novel analytical tools to fuse data across subjects, modalities, and time. Interestingly, once an underlying database of images and anatomical models are collected, statistical (e.g. Bayesian) image analysis tools can be developed to exploit prior empirical data, facilitating each stage of data analysis (Gee et al., 1998; Grenander and Miller, 1998; Ashburner et al., 1997; Le Goualher et al., 1999; Pitiot et al., 2000; Dinov et al., 2000).

1.2 Registration to an Atlas

The need to perform brain-to-brain comparisons has made registration algorithms fundamental to brain image analysis. Images from several subjects, for example, can be analyzed together by first aligning them into a common 3D coordinate space. Motivated by the need to standardize and pool data across subjects, and compare results across laboratories, several registration methods have been developed to align brain mapping data with

an atlas. The simplest registration techniques are linear, removing global differences in brain size. These are covered in detail in Chapter 8. Non-linear approaches, however, can eliminate even the most local size and shape differences that distinguish one brain from another. Transforming individual datasets into the shape of a single reference anatomy, or onto a 3D digital brain atlas, allows subsequent comparison of brain function across individuals (Christensen et al., 1993; Ashburner et al., 1997; Woods et al., 1998). Interestingly, the transformations required to remove individual differences in anatomy are themselves a rich source of morphometric data (Thompson and Toga, 2000), and can be used by probabilistic algorithms for abnormality detection (Section 17.9).

1.2.1 The Talairach System

In the earliest brain atlases, spatial normalization systems were proposed to transform new data to match the space occupied by an atlas. In the Talairach stereotaxic system (Talairach and Tournoux, 1988), piecewise affine transformations were applied to 12 rectangular regions of brain, to re-position the subject's brain in a defined space. Although it was originally designed to assist with localization in surgical studies, the Talairach stereotaxic system rapidly became an international reporting standard for functional activation sites in PET and fMRI studies (Fox et al., 1985, 1988; Friston et al., 1989, 1991).

1.2.2 Digital Templates

While stereotaxic methods provide a common coordinate system to pool activation data for multi-subject comparisons, the accuracy and utility of the underlying atlas depends on the anatomical template itself (Roland and Zilles, 1994). Unfortunately, the original Talairach templates were based on inconsistent, orthogonal sections acquired *post mortem* from a 60 year-old female subject. More recently, digital templates have been created that better reflect the *in vivo* anatomy of subjects in activation studies. This work has also spurred the development of algorithms to automatically align new image data to these templates, rather than relying on the more traditional selection of anatomical landmarks in each individual subject's brain. To facilitate the reporting of brain mapping data in an atlas-based 3D coordinate system, Evans et al. (1994) created a composite MRI dataset from 305 young normal subjects (239 males, 66 females; age: 23.4 ± 4.1 years) whose scans were linearly mapped into stereotaxic space (Talairach and Tournoux, 1988). Aligned scans were intensity normalized, and averaged on a voxel-by-voxel basis (cf. Ashburner et al., 1997). Automated methods were subsequently developed to align new MRI and PET data with these digital templates by optimizing a measure of image similarity, such as 3D cross-correlation

(Collins et al., 1994, 1995), ratio image uniformity (Woods et al., 1992), squared intensity difference (Woods et al., 1993, 1998; Ashburner et al., 1998) or mutual information (Wells et al., 1997). Alignment transformations defined for MRI could also be applied to co-registered PET, SPECT, histologic or anatomical model data, to align them to an atlas. Registration algorithms have therefore made it feasible to automatically map data from a variety of brain imaging modalities into an atlas coordinate space based directly on the Talairach system.

1.3 Deformable Brain Atlases

1.3.1 Atlas to Brain Transformations

Brain structure varies considerably from one individual to another (Steinmetz et al., 1990), so a fixed atlas has obvious limitations. Along with approaches to overlay individual brain maps onto an atlas, new algorithms were rapidly developed to elastically re-shape an atlas to match the anatomy of new individuals. The resulting *deformable brain atlases* can project detailed atlas data into new scans, including maps of cytoarchitecture, biochemistry, functional and vascular territories. Their uses include surgical planning (Warfield et al., 1998; St-Jean et al., 1998), anatomical labeling (Iosifescu et al., 1997) and shape measurement (Subsol, 1995; Thompson et al., 1997; Haller et al., 1997; Csernansky et al., 1998). The shape of the digital atlas is adapted using local warping transformations (dilations, contractions and shearing), producing an *individualized* brain atlas. These transformations allow any segment of the atlas anatomy, however small, to grow, shrink, twist and even rotate, producing a transformation that encodes local differences in topography from one individual to another. The ability to automatically map labeled brain atlases onto individual scans has many applications. Digital anatomic models can be projected into PET data to define regions of interest for quantitative calculations of regional cerebral blood flow (Ingvar et al., 1994; Dinov et al., 2000). Brain structures can also be labeled for hippocampal morphometry in dementia (Haller et al., 1997), for analysis of subcortical structure volumes in schizophrenia (Iosifescu et al., 1997; Csernansky et al., 1998; Gaser et al., 1998), for estimation of structural variation and pathology detection (Collins et al., 1994; Thompson et al., 1997), and for segmentation and classification of multiple sclerosis lesions (Warfield et al., 1995). In view of its broad applications, non-linear registration has a fundamental role in image analysis. When transferring brain image data between scans and atlas templates, clearly the accuracy of the anatomical transformation is paramount. Considerable ingenuity has therefore gone into designing algorithms that use both anatomic and mathematical criteria to reconfigure one anatomy onto another. Because of their importance, these algorithms are reviewed next.

1.4 Warping Algorithms

Non-linear registration approaches are commonly classified into two major types, *intensity-based* and *model-based*, depending on the type of information that drives them (see Toga, 1998, for a review). *Model-driven* algorithms first build explicit geometric models, representing separate, identifiable anatomic elements in each of the scans to be matched. These anatomical systems typically include functionally important surfaces (Szeliski and Lavalle, 1993; Downs et al., 1994; Moshfeghi et al., 1995; Thompson and Toga, 1996; Davatzikos, 1996), curves (Ge et al., 1995; Monga and Benayoun, 1995; Subsol, 1995), and point landmarks (Bookstein, 1989; Amit et al., 1991). Anatomical elements are parameterized and matched with their counterparts in the target scan, and their correspondences guide the volumetric transformation of one brain to another. In our own warping algorithms (Section 17.5; Thompson and Toga, 1996, 1998), higher-level structural information guides the mapping of one brain onto another, and a hierarchy of curve-to-curve and surface-to-surface mappings is set up, guaranteeing the gross anatomical validity of the resulting transform. The algorithms exploit anatomical information to match cortical regions, so that networks of sulci and gyri are individually matched. These strategies are discussed in Section 17.7.

1.4.1 Intensity-Driven Approaches

Model-driven approaches contrast with *intensity-driven approaches*. Intensity-driven approaches aim to match regional intensity patterns in each scan based on mathematical or statistical criteria. Typically, they define a mathematical measure of intensity similarity between the deforming scan and the target. Measures of intensity similarity can include squared differences in pixel intensities (Christensen et al., 1993; Woods et al., 1993, 1998; Ashburner et al., 1997), regional correlation (Bajcsy and Kovacic, 1989; Collins et al., 1995), or mutual information (Kim et al., 1997). Mutual information has proved to be an excellent similarity measure for *cross-modality* registrations, since it assumes only that the *statistical dependence* of the voxel intensities is maximal when the images are geometrically aligned. The intensity similarity measure, often combined with a measure of the structural integrity of the deforming scan, is optimized by adjusting parameters of the deformation field. Images are usually low-pass filtered or represented at multiple scales to make the similarity function smoother with respect to the deformation parameters. This makes it easier to traverse the parameter space and find an optimal match. Nonetheless, algorithms based on intensity patterns alone essentially by-pass information on the internal topology of the brain. Matching of neuroanatomic data in the absence of higher-level structural information presents an extremely difficult pattern recognition

problem. Hybrid algorithms are therefore of particular interest, combining intensity-based and model-based criteria to establish more accurate correspondences (Collins et al., 1996; Grenander and Miller, 1998).

All approaches for intensity-driven image registration define a mathematical measure of intensity similarity between the deforming template and the target image, and optimize it by tuning the parameters of the deformation field. The widely-used *Automated Image Registration* (AIR; Woods et al., 1998) and *Statistical Parametric Mapping* algorithms (Ashburner and Friston, 1999) are examples of this approach. As the cost function (or similarity measure) is optimized, increasingly complex warping fields are expressed in terms of a 3D cosine basis (SPM) or by tuning parameters of 3D polynomials (AIR). In SPM, the target image g($\mathbf{x}$) is approximated by a scaled (by factor w) and spatially deformed version of the individual's image f($\mathbf{x}$). The deformation is constrained to be a linear combination of smooth basis functions:

$$\mathbf{u}(\mathbf{x}) = \left[\sum_j t_{j,1} b_{1,j}(\mathbf{x}), \sum_j t_{j,2} b_{2,j}(\mathbf{x}), \sum_j t_{j,3} b_{3,j}(\mathbf{x})\right] \tag{1.1}$$

where $b_{1,j}$ is the j-th order basis function along axis d at position $\mathbf{x}$. The coefficients $t_{j,d}$ of the deformation field can be assembled, with the intensity scale-factor, into a parameter vector, $\mathbf{p} = [\mathbf{t_x}\ \mathbf{t_y}\ \mathbf{t_z}\ w]$ and their values can be chosen to minimize the least-squares cost function:

$$\sum_i [C(\mathbf{x_i}, \mathbf{p})]^2 = \sum_i [f(\mathbf{y_i}) - wg(\mathbf{x_i})]^2 \tag{1.2}$$

where $\mathbf{y_i}$ is the displaced position of the ith voxel $\mathbf{y_i} = \mathbf{x_i} - \mathbf{u}(\mathbf{x_i})$. To optimize the deformation, note that a small increment $\mathbf{t}$ in the parameter vector will affect the cost function at each voxel i according to the first-order Taylor approximation:

$$C(\mathbf{x_i}, \mathbf{p} + \mathbf{t}) \cong C(\mathbf{x_i}, \mathbf{p}) + t_1[\partial C(\mathbf{x_i}, \mathbf{p})/\partial \mathbf{p_1}] + t_2[\partial C(\mathbf{x_i}, \mathbf{p})/\partial \mathbf{p_2}] + \ldots \tag{1.3}$$

At a global (or local) minimum of the cost function, $\sum_i [C(\mathbf{x_i}, \mathbf{p})]^2$, a linear system $\boldsymbol{At} \cong \boldsymbol{b}$ can be written down and solved for the parameter increment $\mathbf{t}$. Here the matrix elements $A_{rs} = \partial C(\mathbf{x_r}, \mathbf{p})/\partial \mathbf{p_s}$ are computed from the image gradients (using the chain rule), and $\mathbf{t} = [\mathbf{t_1},\ \mathbf{t_2},\ \ldots]^T$ and $\mathbf{b} = [C(\mathbf{x_1}, \mathbf{p}),\ C(\mathbf{x_2}, \mathbf{p}),\ \ldots]^T$. To find the optimal parameters, the deforming image is resampled at each iteration n, and the parameters $\mathbf{p}$ are updated using the Gauss-Newton rule:

$$\mathbf{p^{(n+1)}} = \mathbf{p^{(n)}} - (\mathbf{A^T A})^{-1} \mathbf{A^T b}, \tag{1.4}$$

until the cost function is minimized. Ashburner and Friston (1999) accelerated this scheme by simplifying the large curvature matrix $A^T A$ using known identities for Kronecker tensor products. They also added a Bayesian regularization term that pulls parameter estimates towards their expected values, avoiding unnecessary deformations and accelerating convergence. As in other Bayesian approaches, this covariance term was derived analytically by assuming a Gibbs statistical prior distribution on the deformation energies (cf. Miller et al., 1993). The deformation energy $E(\mathbf{p})$, computed from the transformation parameters, can be transformed into a Gibbs (or Boltzmann) distribution on the expected deformations:

$$P(\mathbf{p}) = (1/Z)e^{-E(\mathbf{p})}, \tag{1.5}$$

Here Z is the partition function that normalizes the distribution. In the SPM approach, the covariance matrix of the deformation parameters is expanded in terms of the eigenfunctions of the governing operator (here the DCT basis functions), and used to add a Bayesian prior term that pulls the mapping away from unrealistic deformations.

1.4.2 Bayesian Methods

In a related Bayesian approach (Gee and Le Briquer, 1997), deformation mappings are constrained to lie in the space of normal anatomical variability, which is estimated empirically from a set of inter-subject mappings. If the Cartesian components of each deformation field are stacked into a high-dimensional vector, $\mathbf{u^{(i)}}(\mathbf{x})$, the covariance matrix of the mappings will be singular because the number of observed mappings is small compared with their dimensionality. Gee and Le Briquer (1997) addressed this problem by deriving a new orthogonal basis on the deformation space by Gram-Schmidt orthogonalization. Re-expressed in new basis, the mean and covariance matrix of the deformation coefficients are used to derive a Gaussian prior on the deformation space. A linear system is then solved for the mapping that optimizes a combination of least-squares intensity similarity and prior probability, as quantified by the empirical distribution. As the principal modes of deformation are computed in advance, the resulting mappings are computed rapidly, guided by empirical knowledge on brain shape variability.

1.4.3 Polynomial Mappings

The widely-used *Automated Image Registration* (Woods et al., 1998) package also uses the least-squares cost function (with intensity scaling) for nonlinear registration, but expresses each component u^k of the deformation field as a polynomial (of up to 12th order) in the coordinates of the target image. Again, parameter adjustments $\boldsymbol{t}$ are computed iteratively, this time by solving a related, second-order linear system, $\boldsymbol{Ht} = -\boldsymbol{d}$. Here $\boldsymbol{H}$ is the

symmetric Hessian matrix of second derivatives of the cost function with respect to the spatial transformation parameters, and $\boldsymbol{d}$ is the gradient of cost function. This formulation results from approximating the cost function by its tangential quadratic form at the current parameters $\boldsymbol{p}$:

$$C(\mathbf{x_i}, \mathbf{p} + \mathbf{t}) \cong C(\mathbf{x_i}, \mathbf{p}) + < \nabla C(\mathbf{x_i}, \mathbf{p}), \mathbf{t} > + \mathbf{t^T}.\mathbf{H}[C(\mathbf{x_i}, \mathbf{p})].\mathbf{t} + \ldots \quad (1.6)$$

This quadratic approximation is minimized when its gradient is the zero vector:

$$0 = \nabla C(\mathbf{x_i}, \mathbf{p}) + \mathbf{H}[C(\mathbf{x_i}, \mathbf{p})].\mathbf{t}, \quad (1.7)$$

so the method for incrementing parameters is given by:

$$\mathbf{t} = -\mathbf{H}[C(\mathbf{x_i}, \mathbf{p})]\nabla C(\mathbf{x_i}, \mathbf{p}). \quad (1.8)$$

Anatomical validation experiments (Woods et al., 1998) showed considerable improvements in registration accuracy with polynomials up to 6th order, when aligning major cortical landmarks across subjects.

1.4.4 Continuum-Mechanical Transformations

Both SPM and AIR express deformation fields using global deformation functions, and the complexity of the mappings is generally not increased beyond 8x8x8 basis functions or 8th order polynomial mappings. Higher-order transformations are required for applications that exactly reconfigure one anatomy onto another, such as tracking intraoperative brain deformation, or performing a locally accurate image segmentation. Physical continuum models, for example (Fig. 2), allow extremely flexible deformations, potentially with as many degrees of freedom as there are voxels in the image. These approaches consider the deforming image to be embedded in a 3D deformable medium, which can be either an elastic material or a viscous fluid. The medium is subjected to distributed internal forces, which reconfigure the image to match the target. These forces can also be based on the local intensity patterns in the datasets, to match image regions of similar intensity.

1.4.5 Navier-Stokes Equilibrium Equations

In elastic media, the displacement field $\mathbf{u}(\mathbf{x})$ resulting from internal deformation forces $\mathbf{F}(\mathbf{x})$ (called *body forces*) obeys the Navier-Stokes equilibrium equations for linear elasticity:

$$\mu\nabla^2\mathbf{u} + (\lambda + \mu)\nabla(\nabla^T \bullet \mathbf{u}(\mathbf{x})) + \mathbf{F}(\mathbf{x} - \mathbf{u}(\mathbf{x})) = \mathbf{0}, \forall \mathbf{x} \in \mathbf{R}, \quad (1.9)$$

Here $\mathbf{R}$ is a discrete lattice representation of the scan to be transformed, $\nabla^T \bullet \mathbf{u}(\mathbf{x}) = \sum_j \partial u_j / \partial x_j$ is the divergence, or cubical dilation of the medium, ∇^2 is the Laplacian operator, and Lamé's coefficients λ and μ refer to the elastic properties of the medium (see Fig. 2). Body forces, designed to match regions in each dataset with high intensity similarity, can be derived from the gradient of a cost function, such as intensity correlation. In Bajcsy and Kovacic (1989), intensity neighborhoods to be correlated in each scan were projected onto a truncated 3D Hermite polynomial basis to enhance the response of edge features and accelerate computation. More complex local operators can also be designed to identify candidates for regional matches in the target image (Amit, 1997). The elasticity equations can then be solved in a variety of ways, using finite differences (Broit, 1981), multiresolution/multigrid schemes (Bajcsy and Kovacic, 1989), finite elements (Gee et al., 1993) or spectral methods (Miller et al., 1993).

1.4.6 Viscous Fluid Approaches

More recently, Christensen et al. (1993, 1995, 1996) proposed a viscous-fluid based warping transform, motivated by capturing non-linear topological behavior and large image deformations (see also Dupuis et al., 1998). Additional continuum-mechanical constraints use the properties of fluids to guarantee the topological integrity of the deformed template. Similar to SPM, a low-order deformation is computed first in terms of an approximation series of eigenfunctions of the linear elasticity operator $\mu\nabla^2 + (\lambda + \mu)\nabla(\nabla^T \bullet)$. This basis function representation of the deformation is analogous to the discrete cosine basis used in SPM (which corresponds to the Laplacian operator ∇^2. The elastic eigenfunctions penalize extreme dilation and compression of the deformed image (Fig. 2), via an additional gradient-of-the-divergence term $\nabla(\nabla^T \bullet)$ not present in the Laplacian formulation. Basis coefficients are determined by gradient descent on a cost functional (10) that penalizes squared intensity mismatch between the deforming template $T(\mathbf{x} - \mathbf{u}(\mathbf{x},t))$ and target $\mathbf{S}(\mathbf{x})$:

$$C(\mathbf{T}(\mathbf{x}), \mathbf{S}(\mathbf{x}), \mathbf{u}) = (1/2) \int_\Omega |T(\mathbf{x} - \mathbf{u}(\mathbf{x},t)) - \mathbf{S}(\mathbf{x})|^2 \mathbf{dx} \tag{1.10}$$

By contrast with SPM and AIR, stochastic gradient descent is used to find the optimal warping field parameters according to:

$$d\mu_{i,j,r}(t) = -(1/2)[(\partial H(\mathbf{u}(t))/d\mu_{i,j,r})]dt + dw_{i,j,r}(t) \tag{1.11}$$

Here $\mu_{i,j,r}$ is the expansion coefficient set for the deformation field in terms of the eigenbasis $\{\mathbf{e}_{i,j,r}\}$ for the linear elasticity operator, $H(\mathbf{u}(t))$ is the combined measure of intensity mismatch and deformation severity, and

$dw_{i,j,r}(t)$ is a Wiener process that allows provisional parameter estimates to jump out of local minima. At the expense of added computation time, stochastic sampling allows globally optimal image matches to be estimated. Finally, a viscous deformation stage allows large-distance, non-linear fluid evolution of the neuroanatomic template. With the introduction of concepts such as deformation velocity and a Eulerian reference frame, the energetics of the deformed medium are hypothesized to be relaxed in a highly viscous fluid. The driving force, which deforms the anatomic template, is defined as the variation of the cost functional with respect to the displacement field:

$$F(\mathbf{x}-\mathbf{u}(\mathbf{x},t)) = -(T(\mathbf{x}-\mathbf{u}(\mathbf{x},t)) - \mathbf{S}(\mathbf{x}))\nabla T|_{\mathbf{x}-\mathbf{u}(\mathbf{x},t)} \tag{1.12}$$

$$\mu\nabla^2\mathbf{v}(\mathbf{x},t) + (\lambda+\mu)\nabla(\nabla^T \bullet \mathbf{v}(\mathbf{x},t)) + \mathbf{F}(\mathbf{x},\mathbf{u}(\mathbf{x},t)) = \mathbf{0} \tag{1.13}$$

$$\partial\mathbf{u}(\mathbf{x},t)/\partial t = \mathbf{v}(\mathbf{x},t) - \nabla\mathbf{u}(\mathbf{x},t)\mathbf{v}(\mathbf{x},t) \tag{1.14}$$

The deformation velocity (13) is governed by the creeping flow momentum equation for a Newtonian fluid and the conventional displacement field in a Lagrangian reference system (14) is connected to a Eulerian velocity field by the relation of material differentiation. Experimental results were excellent (Christensen et al., 1996).

1.4.7 Acceleration with Fast Filters

Vast numbers of parameters are required to represent complex deformation fields. In early implementations, deformable registration of a 128^3 MRI atlas to a patient took 9.5 and 13 hours for elastic and fluid transforms, respectively, on a 128x64 DECmpp1200Sx/Model 200 MASPAR (Massively Parallel Mesh-Connected Supercomputer). This spurred work to modify the algorithm to individualize atlases on standard single-processor workstations (Thirion, 1995; Bro-Nielsen and Gramkow, 1996; Freeborough and Fox, 1998).

Bro-Nielsen and Gramkow (1996) used the eigenfunctions of the Navier-Stokes differential operator $L = \mu\nabla^2 + (\lambda+\mu)\nabla(\nabla^T\bullet)$, which governs the atlas deformations, to derive a Green's function solution $\mathbf{u}^*(\mathbf{x}) = \mathbf{G}(\mathbf{x})$ to the impulse response equation $L\mathbf{u}^*(\mathbf{x}) = \delta(\mathbf{x}-\mathbf{x_0})$. This speeds up the core registration step by a factor of 1000. The solution to the full PDE $L\mathbf{u}(\mathbf{x}) = -\mathbf{F}(\mathbf{x})$ was approximated as a rapid filtering operation on the 3D arrays representing body force components:

$$\mathbf{u}(\mathbf{x}) = -\int \mathbf{G}(\mathbf{x}-\mathbf{r}).\mathbf{F}(\mathbf{r})\mathbf{dr} = -(\mathbf{G}*\mathbf{F})(\mathbf{x}), \tag{1.15}$$

where $\mathbf{G}*$ represents convolution with the impulse response filter. As noted in (Gramkow and Bro-Nielsen, 1997), a recent fast, 'demons-based' warping algorithm (Thirion, 1995; Dawant et al., 1998; Cachier et al., 1999)

calculates the atlas flow velocity by regularizing the force field driving the template with a Gaussian filter (cf. Collins et al., 1994). Since this filter may be regarded as a separable approximation to the continuum-mechanical filters derived above (Nielsen et al., 1994), interest has focused on deriving additional separable (and therefore computationally fast) filters to create subject-specific brain atlases and rapidly label new images (Gramkow, 1996; Lester et al., 1999). Ultimately, filtering the driving force, as well as the deformation field (or its increments, (eqn. 13)) are central to high-dimensional non-linear registration. With this in mind, Cachier et al. (1999) developed an *a posteriori* filter-weighting approach that attenuates the weight of the driving force at positions where it leads to a poorer match. Fast multi-grid solvers have also accelerated systems for atlas-based segmentation and labeling (Dengler and Schmidt, 1988; Bajcsy and Kovacic (1989); Collins et al., 1994, 1995; Gee et al., 1993,1995; Schormann et al., 1996). Some of these now have sufficient speed for real-time surgical guidance applications (Warfield et al., 1998).

1.4.8 Neural Network Implementations

There is an interesting mathematical connection between continuum-mechanical PDEs and neural nets, which has been exploited to generate fast algorithms for brain image registration. Neural network algorithms use the fact that the simplest set of anatomic features that can guide the mapping of one brain to another is a set of point landmarks. Point correspondences can be extended to produce a deformation field for the full volume in a variety of ways, each consistent with the displacements assigned at the point landmarks. The ambiguity is resolved by requiring the deformation field to be the one that minimizes a specific *regularizing functional* (Tikhonov and Arsenin, 1977). This measures the roughness or irregularity of the deformation field, calculated from its spatial derivatives. *Thin-plate splines* (Grimson, 1981; Bookstein, 1989), *membrane splines* (Amit et al., 1991; Gee et al., 1993), *elastic body splines* (Miller et al., 1993; Davis et al., 1997), and *div-curl splines* (Suter, 1994) are functions which minimize the following distortion measures (in 2D):

$$J_{thin-plate}(\mathbf{u}) = \int_{R^2} [(\partial_{11}\mathbf{u})^2 + 2(\partial_{12}\mathbf{u})^2 + (\partial_{22}\mathbf{u})^2] dx_1 dx_2 \quad (1.16)$$

$$J_{membrane}(\mathbf{u}) = \int_{R^2} [(\partial_1 u_1)^2 + (\partial_1 u_2)^2 + (\partial_2 u_1)^2 + (\partial_2 u_2)^2] dx_1 dx_2 \quad (1.17)$$

$$J_{elastic}(\mathbf{u}) = \int_{R^2} \sum_{i=1}^{2} \sum_{j=1}^{2} [(\lambda/2)(\partial_i u_i)(\partial_j u_j) + (\mu/4)(\partial_i u_j)(\partial_j u_i)] dx_1 dx_2 \quad (1.18)$$

$$J_{div-curl}(\mathbf{u}) = \int_{R^2} [(\lambda \parallel \nabla \mathbf{DIVu} \parallel^2 + (\mu \parallel \nabla \mathbf{CURLu} \parallel^2] dx_1 dx_2 \quad (1.19)$$

where $\partial_{ij}\mathbf{u} = \partial^2\mathbf{u}/\partial x_i \partial x_j$ [1].

Once a type of spline is chosen for warping, a formula can be used which specifies how to interpolate the displacement field from a set of points $\{\mathbf{x_i}\}$ to the surrounding 2D plane or 3D volume:

$$\mathbf{u}(\mathbf{x}) = p_{m-1}(\mathbf{x}) + \sum_i c_i G(\mathbf{x} - \mathbf{x_0}) \tag{1.20}$$

Here $p_{m-1}(\mathbf{x})$ is a polynomial of total degree $m-1$, where m is the order of derivative used in the regularizer, and G is a *radial basis function* (RBF) or *Green's function* whose form depends on the type of spline being used (Joshi et al., 1995; Davis et al., 1997). Choices of $r^2 \ln r$ and r correspond to the *thin-plate spline* in 2D and 3D, with r^3 for the *3D volume spline* (Davis et al., 1997), and the 3x3 matrix $[(r^2\mathbf{I} - \mathbf{3xx^T}]r$ for the *3D elastic body spline* (Davis et al., 1997). Substitution of the point correspondences into this formula results in linear system that can be solved for the deformation field (Fig. 2; Thompson and Toga, 1998). Neural network approaches exploit this by using correspondences at known landmarks as a *training set* to learn a multivariate function. This function maps positions in the image (input) to the desired displacement field at that point (output). Intriguingly, the hidden units in the neural net are directly analogous to Green's functions, or convolution filters, in the continuum-mechanical matching approach (Joshi et al., 1995; Bro-Nielsen and Gramkow, 1996). They are also directly analogous to Watson-Nadaraya *kernel estimators*, or Parzen windows, in non-parametric regression methods (Parzen, 1962). By converting the above linear system into a neural network architecture, the k deformation field components are the output values of the neural net:

$$u^k(\mathbf{x}) = \sum_{m=1}^{M} a_m \pi_m(\mathbf{x}) + \sum_{i=1}^{N} w_{ik} G_i(\mathbf{x} - \mathbf{x_i}) \tag{1.21}$$

Here the G_i are N separate hidden unit neurons with receptive fields centered at x_i, $\sum_{m=1}^{M} a_m \pi_m(\mathbf{x})$ is a polynomial whose terms are hidden units and whose coefficients a_m are also learned from the training set, and w_{ik} are synaptic weights (Fig. 2). The synaptic weights are determined by solving a linear system obtained by substituting the training data into this equation. If landmarks are available to constrain the mapping, the function centers x_i may be initialized at the landmark positions, otherwise hidden units can

[1] Just like the continuum-mechanical warps defined earlier, the warping fields generated by splines satisfy partial differential equations of the form $L\mathbf{u}(\mathbf{x}) = -\mathbf{F}(\mathbf{x})$, where $\mathbf{u}(\mathbf{x})$ is fixed at the specified points, $\mathbf{F}(\mathbf{x})$ plays the role of a body force term, and L is the biharmonic differential operator (∇^4 for the thin-plate spline, the Laplacian operator ∇^2 for the membrane spline, and the Cauchy-Navier operator $\mu\nabla^2 + (\lambda + \mu)\nabla(\nabla^T \bullet)$ for the elastic body spline.

initially be randomly placed across the image (Davis et al., 1996). Network weights (the coordinate transformation parameters) and the RBF center locations are successively tuned to optimize an intensity-based functional (normalized correlation) that measures the quality of the match. The network is trained (i.e., the parameters of the warping field are determined) by evaluating the gradient of the normalized correlation with respect to the network parameters, and optimizing their values by gradient descent. Results matching 3D brain image pairs were impressive (Davis et al., 1996). For further discussion of the close relationship between continuum-mechanical PDEs, statistical regression and neural nets, see Ripley et al. (1996).

1.5 Model-Driven Deformable Atlases

Deformable atlases, driven by nonlinear registration algorithms, can be regarded as a technique for automatically labeling brain data. If they are accurate, they can find structures in new images, for subsequent manipulation and analysis. However, the extreme difficulty of finding structures in new patients based on intensity criteria alone has led several groups to develop model-driven deformable atlases (Thompson and Toga, 1997; Toga and Thompson, 1997). Anatomical models provide an explicit geometry for individual structures in each scan, such as landmark points, curves or surfaces. To compute the anatomical differences between two subjects, or between a subject and an atlas, corresponding models can be matched and a vector field computed to reconfigure one set of models into the shape of the other. Because the digital models reside in the same stereotaxic space as the atlas data, their vector coordinates are amenable to digital transformation, as well as geometric and statistical measurement (Thompson et al., 1996). The underlying 3D coordinate system is central to all atlas systems, since it supports the linkage of structure models and associated image data with spatially-indexed neuroanatomic labels, preserving spatial information and adding anatomical knowledge.

1.5.1 Anatomical Modeling

In the following sections, we show how anatomical models can be used to create probabilistic brain atlases and disease-specific templates. Statistical averaging of models provides a means to analyze brain morphometry, localizing disease-specific differences with statistical and visual power. We first describe how models can drive deformable atlases, measuring patient-specific differences in considerable detail. When deforming an atlas to match a patient's anatomy, mesh-based models of anatomic systems guide the mapping of one brain to another. Anatomically-driven algorithms guarantee biological as well as computational validity, generating meaningful object-to-object correspondences, especially at the cortex.

1.5.2 Parametric Meshes

Since much of the functional territory of the human cortex is buried in cortical folds or *sulci*, considerable attention has focused on building a generic structure to model them (Fig. 3; Thompson and Toga, 1996). In our own surface parameterization approach (Thompson et al., 1996, 1997, 2000), we model anatomy using systems of surface meshes, in which the individual meshes are parametric. These surfaces are 3D sheets that divide and join at curved junctions to form a connected network of models. With the help of these meshes, each patient's anatomy is represented in sufficient detail to be sensitive to subtle disease-specific differences. The parametric grid imposed on each element of the anatomy provides a computational structure that supports (1) measurement of geometric shape parameters (e.g., curvature, area, complexity, fractal dimension; see Thompson et al., 1996 for details); (2) combination of models across subjects to produce average models and statistical maps (Thompson et al., 1996, 1997; Narr et al., 2000); and (3) measurement of gyral pattern differences, by discretizing partial differential equations that compute flows to match cortical surfaces (Thompson et al., 1997, 2000; Sowell et al., 2000; Narr et al., 2000; Blanton et al., 2000).

To identify and analyze patterns of altered structure in disease, large systems of anatomical models can be stored in a population-based atlas. In these morphometric atlases, separate surfaces model the deep internal trajectories of features such as the parieto-occipital sulcus, the anterior and posterior calcarine sulcus, the Sylvian fissure, and the cingulate, marginal and supracallosal sulci in both brain hemispheres. Additional gyral boundaries are represented by parameterized curves lying in the cortical surface. The ventricular system is modeled as a closed system of 14 connected surface elements whose junctions reflect the cytoarchitectonic boundaries of the adjacent tissue (Fig. 5; Thompson and Toga, 1998). Information on the meshes' spatial relations, including their surface topology (*closed* or *open*), anatomical names, mutual connections, directions of parameterization, and common 3D junctions and boundaries is stored in a hierarchical graph structure. This ensures the continuity of displacement vector fields defined at mesh junctions.

1.5.3 Automated Parameterization

A major goal in brain mapping is to find surfaces of brain structures automatically. If an identical regular grid structure can be imposed on anatomic surfaces from different subjects (Fig. 4), the explicit geometry can be exploited to compute shape measures and correspondence maps that associate anatomic points in different subjects. Accurate labeling of anatomy, especially at the cortex, requires detailed case-by-case rules that are difficult to formulate computationally (Ono et al., 1990; Thompson et al., 1997; Sowell

et al., 2000). Current approaches for automated parameterization of brain structures fall into two major categories: (1) *deformable templates*, and (2) *voxel coding*. Deformable templates are covered in detail in Chapters 3 and 6. Briefly, the shape of a prototype model, such as a parametric curve or surface, is tuned until a measure of fit is optimized, suggesting that the target object has been found in the image. Probabilistic brain atlases can assist these algorithms, in that (1) stored information on empirical shape variability can guide deformable templates (Cootes et al., 1995), and (2) a validation test-bed of existing models can be used to optimize algorithm parameters. To investigate this, we recently developed a Bayesian approach to identify the *corpus callosum* in each image in an MRI database (Pitiot et al., 2000; cf. Staib and Duncan, 1992). The shape of a deformable curve (Fig. 6, *panel 7*) is progressively tuned to optimize a mathematical criterion measuring how likely it is that it has found the corpus callosum. The measure includes terms that reward contours based on their agreement with a diffused edge map (*panels 7-9*), their geometric regularity, and their statistical abnormality when compared with a distribution of normal shapes. As described by Cootes et al. (1995) and Jain et al. (1996), a preference for specific shapes is expressed by specifying a probability distribution on the parameters of the deformation function. Since the best algorithm parameters may depend on the image noise (e.g. the size of the connectivity filter for edge suppression), optimal values were determined empirically from simulations on a database of 104 brain images. The performance of the algorithm is shown in Fig. 6. As we shall see later, by averaging *corpus callosum* contours derived from an image database, structural abnormalities associated with Alzheimer's Disease, schizophrenia and fetal alcohol syndrome can be identified and visualized (Thompson et al., 1998; Narr et al., 2000; Sowell et al., 2000).

1.5.4 Voxel Coding

Voxel-coding (Zhou et al., 1999, 2000) provides a fundamentally different approach for automated structure parameterization. Rather than starting with an a priori geometric model of the target structure, low-level image operations, such as erosion, dilation, are applied at a voxel-by-voxel level to gradually build up parametric surface grids. The term voxel-coding derives from the procedure of repeatedly assigning numerical codes to voxels. These codes are used to sort voxels efficiently, find the shortest paths in 2D and 3D, and accelerate mathematical morphology operations such as erosion and skeletonization. This by-passes the complex mathematics of deformable templates. The resulting approach is computationally fast, and extracts parameterized models of multiple surfaces in a 3D image at once. A related approach can be used to extract sets of curved lines representing the superficial

sulcal pattern (Lohmann, 1999). In view of the interest in deformable surfaces to parameterize deep sulcal anatomy (Vaillant et al., 1997; Le Goualher et al., 1999), we recently developed a voxel-coding approach to address the problem (Zhou et al., 1999; Fig. 7). In this approach, a supervised classifier is used derive a binary map of gray matter and cerebro-spinal fluid (CSF) which contains the sulci. In each 2D slice a fast algorithm propagates a distance field from the image exterior to the sulcal beds. Local maxima of the distance function are identified, and from these points shortest voxel paths to the exterior are identified. These paths are adjusted to a medial course by reference to a second distance field propagated from the sulcal banks. A third distance field is then traversed to establish local object connectivity between slices, and the resulting voxel set is uniformly re-parameterized and triangulated (Thompson et al., 1996). Topological issues of ambiguous or multiple connectivity, e.g. when objects merge or divide, are addressed by introducing a simplified skeleton extracted from a region that defines inter-slice differences (Zhou et al., 2000). Automatically extracted models agreed well with manual-derived data, and their local accuracy was mapped using an adaptive approach based on Hotelling's T^2 random fields to encode patterns of manual error (Thompson et al., 1997; Zhou et al., 1999). Work is underway to integrate this approach into an automated image analysis pipeline for analyzing disease-related patterns of anatomy.

1.5.5 Model-Based Deformable Atlases

Parametric models can also be used to measure cross-subject anatomical differences, by computing the amount of deformation required to reconfigure one anatomy into the shape of another. We recently developed a surface-based 3D image warping algorithm, which matches complex anatomical surface boundaries when driving one brain into the shape of another. Specialized approaches are used to match gyral patterns of the cortex across subjects (see Section 17.7) constraining the anatomical transformation with anatomical landmark points, curves, surfaces, and even curves within surfaces. For each surface mesh M_l in a pair of scans A_p and A_q we define a 3D displacement field:

$$\mathbf{W}_{\mathbf{l}}^{\mathbf{pq}}[\mathbf{r}_{\mathbf{l}}^{\mathbf{p}}(u,v)] = \mathbf{r}_{\mathbf{l}}^{\mathbf{q}}(u,v) - \mathbf{r}_{\mathbf{l}}^{\mathbf{p}}(u,v) \tag{1.22}$$

carrying each surface point $\mathbf{r}_{\mathbf{l}}^{\mathbf{p}}(u,v)$ in A_p into structural correspondence with $\mathbf{r}_{\mathbf{l}}^{\mathbf{q}}(u,v)$, the point in the target mesh parameterized by rectangular coordinates (u,v). This family of high-resolution transformations, applied to individual meshes in a connected system deep inside the brain, elastically transforms elements of the surface system in one 3D image to their counterparts in the target scan. Weighted linear combinations of radial functions, describing the influence of deforming surfaces on points in their vicinity, ex-

tend the surface-based deformation to the whole brain volume (see Fig. 8). Recent extensions of the core algorithm to include continuum-mechanical, and other filter-based models of deformation (Thompson and Toga, 1998; cf. Joshi et al., 1995; Davatzikos, 1996; Schiemann and Hhne, 1997; Gabrani and Tretiak, 1999) have yielded similar encouraging results. Fig. 8 shows how the algorithm performs on post mortem cryosectioned data.

1.6 Probabilistic Atlases and Model-Based Morphometry

1.6.1 Anatomical Modeling

Many morphometric studies focus on identifying systematic alterations in anatomy in a variety of diseases. These studies are complicated by the substantial overlap between measures of normal and diseased anatomy, which makes group-specific patterns hard to discern. However, recent studies suggest that disease-specific variants may be easier to localize by creating average models of anatomy, rather than deriving volumetric descriptors (Thompson et al., 1997; Narr et al., 2000a,b; Sowell et al., 2000a,b; Blanton et al., 2000).

In response to these challenges, *probabilistic atlases* are research tools that retain information on anatomic and functional variability (Mazziotta et al., 1995; Thompson et al., 1997, 2000). A probabilistic atlas solves many of the limitations of a fixed atlas in representing highly variable anatomy. As the subject database increases in size and content, the digital form of the atlas allows efficient statistical comparisons of individuals or groups. In addition, the population that an atlas represents can be stratified into subpopulations to represent specific disease types, and subsequently by age, gender, handedness, or genetic factors.

1.6.2 Parametric Mesh Models

Parametric meshes (Thompson et al., 1996a,b) offer a means to create average models of anatomy. Once anatomical data is transformed to a standardized coordinate space, such as the Talairach space, a computational grid structure can be imposed on anatomical surface boundaries. These mesh models represent boundary point locations in stereotaxic coordinates (Fig. 4). Averaging of corresponding grid points across subjects results in an average surface model for each structure. At the same time, knowledge of each subject's deviations from the anatomical group average can be retained as a vector displacement map (Fig. 4). After storing these maps from large numbers of subjects, local biases in the magnitude and direction of anatomic variability can be displayed as a map.

1.6.3 3D Maps of Variability and Asymmetry

Fig. 9 shows average maps of the lateral ventricles, again from Alzheimer's disease and matched elderly normal populations. In these maps, the color shows the root mean square magnitude of the displacement vectors that map individuals to the group mean. Separate maps are displayed for elderly normals (mean age: 72.9±5.6 yrs.), and Alzheimer's disease (AD) patients matched for age, gender, handedness (all right-handed) and educational level (age: 71.9±10.7 yrs.; mean Mini-Mental State Exam score: 19.7±5.7, out of 30). Clearly, the ventricles are significantly enlarged in dementia. By contrast with conventional volumetric approaches, which indicate that the ventricles are enlarged overall in AD, the region of greatest disease-related enlargement is clearly localized to the occipital horn. A second feature observable from the average anatomical models (Fig. 9) is that consistent patterns of brain asymmetry can be mapped, despite wide variations in asymmetry in individual subjects. This is an example of an effect that becomes clear after group averaging of anatomy, and is not universally apparent in individual subjects. On average, the occipital horn extends 5.1 mm more posteriorly on the left than the right. Anatomical averaging can also be cross-validated with a traditional volumetric approach. In this subject group, occipital horns of the ventricles were on average 17.1% larger on the left in the normal group (4070.1±479.9 mm^3) than on the right (3475.3±334.0 mm^3; $p \prec 0.05$), but no significant asymmetry was found for the superior horns ($left$: 8658.0±976.7 mm^3; $right$: 8086.4±1068.2 mm^3; $p \prec 0.19$) or for the inferior horns ($left$: 620.6±102.6 mm^3; $right$: 573.7±85.2 mm^3; $p \prec 0.37$). The asymmetry is clearly localized in the 3D group average anatomic representations.

1.6.4 Alzheimer's Disease

We also tested the ability of anatomical averaging to identify disease-specific patterns in clinical populations. First, we applied the approach to detect pre-clinical hippocampal atrophy in patients with minimal cognitive impairment (Kwong et al., 1999; Mega et al., 2000). To identify more focal effects, we attempted to identify regionally selective patterns of callosal change in patient groups with Alzheimer's disease and schizophrenia (Thompson et al., 1998; Narr et al., 2000). The mid-sagittal callosum was first partitioned into 5 sectors (Fig. 10; Duara et al., 1991; Larsen et al., 1992) that roughly segregate callosal fibers from distinct cortical regions. In AD, focal fiber loss was expected at the callosal isthmus (sector 2), which carries fibers that selectively innervate the temporo-parietal regions that show early neuronal loss and perfusion deficits (Brun and Englund, 1981). Consistent with this hypothesis, a significant area reduction at the isthmus was found, reflecting a dramatic 24.5% decrease from 98.0±8.6 mm^2 in controls to 74.0±5.3

mm^2 in AD ($p \prec 0.025$). Terminal sectors (1 and 5) were not significantly atrophied, and the central midbody sector showed only a trend toward significance (16.6% mean area loss; $p \prec 0.1$), due to substantial inter-group overlap. Average boundary representations, however, localized these findings directly. At the *isthmus*, average models in AD showed a pronounced shape inflection at stereotaxic location (0.0,-25.0,19.0) (see Fig. 10).

1.6.5 Gender in Schizophrenia

Different shape alterations were observed in schizophrenia (Narr et al., 1999; Fig. 11). A significant bowing effect was observed, reflecting enlargement of the underlying superior and posterior horns of the lateral ventricles. By creating separate average models for male and female patients, significant gender effects also emerged (Fig. 11(a)-(d)). The greater bowing effect in male than female patients was confirmed by multivariate analysis of variance, and is highlighted in the average anatomic templates. As emphasized by this example, even if no sex difference is present in normal callosal morphology (see Thompson et al., 1999, for a review of this controversy), this does not preclude sex effects from interacting with morphometric abnormalities in diseased populations. In schizophrenia, there is typically a later age of onset in female schizophrenics, and hereditary factors may be unevenly distributed between the sexes (De Lisi et al., 1989; Waddington, 1993; Colombo et al., 1993). Stratification of probabilistic atlases by gender and other genetic factors provides a computationally fast way to visualize these effects and relate them to epidemiologic data (Mazziotta et al., 1995; Mega et al., 1998; Zoumalan et al., 1999; Blanton et al., 1999; Le Goualher et al., 1999). The capacity to resolve group features in a population-based atlas can also assist in studies of disease-specific cortical organization (Thompson et al., 1997; Mega et al., 1998; Zoumalan et al., 1999; Narr et al., 1998, 2000).

1.7 Cortical Modeling and Analysis

The cortex presents unique challenges in brain mapping. Extreme variations in gyral patterns make it hard to (1) identify homologous cortical regions across subjects, (2) pool cortically-derived imaging data across subjects, and (3) distinguish normal from abnormal cortical structure. The cortex is also severely affected in disorders such as Alzheimer's disease, Pick's disease and other dementias, by tumor growth, and in cases of epilepsy, cortical dysplasias, and schizophrenia. A major challenge in investigations of disease is to determine (1) whether cortical organization is altered, and if so, which cortical systems are implicated, and (2) whether normal features of cortical organization are lost, such as sulcal pattern asymmetries. This requires methods to create a well-resolved average model of the cortex specific for a diseased group, and a statistical framework to compare individual

and group average models with normative data.

1.7.1 Cortical Matching

Cortical anatomy can be compared, between any pair of subjects, by computing the warped mapping that elastically transforms one cortex into the shape of the other. Due to variations in gyral patterning, cortical variability will be severely underestimated unless elements of the gyral pattern are matched from one subject to another. This matching is also required for cortical averaging; otherwise, corresponding gyral features will not be averaged together. Transformations can therefore be developed that match large networks of gyral and sulcal features with their counterparts in the target brain (Thompson and Toga, 1996, 1997; Davatzikos, 1996; Van Essen et al., 1997; Fischl et al., 1999). Differences in cortical organization prevent exact gyrus-by-gyrus matching of one cortex with another. In our approach, we aim to match 38 elements of the gyral pattern, including the major features that are consistent in their incidence and topology across subjects (Thompson et al., 1997; Sowell et al., 2000; cf. Ono et al., 1990; Leonard et al., 1996; Kennedy et al., 1998).

To find good matches among cortical regions we perform the matching process in the cortical surface's parametric space, which permits more tractable mathematics (Fig. 12,13). This vector flow field in the parametric space indirectly specifies a correspondence field in 3D, which drives one cortical surface into the shape of another. This mapping not only matches overall cortical geometry, but matches the entire network of the 38 landmark curves with their counterparts in the target brain, and thus is a valid encoding of cortical variation.

1.7.2 Spherical, Planar Maps of Cortex

Several simpler maps of the cortex are made to help calculate the transformation. Cortical models are often created by deforming a spherical mesh into the shape of the cortex (MacDonald et al., 1993; Davatzikos, 1996). Any point on the cortex then maps to exactly one point on the sphere, and a *spherical map* of the cortex can be made which indexes sulcal landmarks in the normally folded brain surface. These spherical locations, indexed by two parameters, can also be mapped to a plane (Fig. 12,13). A flow field is then calculated that elastically warps one flat map onto the other (or equivalently, one spherical map to the other). On the sphere, the parameter shift function $\mathbf{u}(\mathbf{r}) : \mathbf{\Omega} \rightarrow \mathbf{\Omega}$ is given by the solution $F_{pq} : \mathbf{r} \rightarrow \mathbf{r} - \mathbf{u}(\mathbf{r})$ to a curve-driven warp in the spherical parametric space $\Omega = [0, 2\pi) \times [0, \pi)$ of the cortex (Thompson et al., 1997; cf. Bakircioglu et al., 1999). For points $\mathbf{r} = (r, s)$ in the parameter space, a system of simultaneous partial differential equations can be written for the flow field $\mathbf{u}(\mathbf{r})$:

$$L^{\ddagger}\mathbf{u}(\mathbf{r}) + \mathbf{F}(\mathbf{r} - \mathbf{u}(\mathbf{r})) = \mathbf{0}, \forall \mathbf{r} \in \mathbf{\Omega}, \quad with \quad \mathbf{u}(\mathbf{r}) = \mathbf{u_0}(\mathbf{r}), \forall \mathbf{r} \in \mathbf{M_0} \cup \mathbf{M_1}, \tag{1.23}$$

Here M_0, M_1 are sets of points and (sulcal or gyral) curves where displacement vectors $\mathbf{u}(\mathbf{r}) = \mathbf{u_0}(\mathbf{r})$ matching corresponding anatomy across subjects are known. The flow behavior is modeled using equations derived from continuum mechanics, and these equations are governed by the Cauchy-Navier differential operator $L = \mu\nabla^2 + (\lambda+\mu)\nabla(\nabla^T \bullet)$ with body force $\mathbf{F}$ (Thompson et al., 1996, 1998, 1999; Grenander and Miller, 1998). The only difference is that $L^{\ddagger}$ is the *covariant* form of the differential operator L, for reasons explained next.

1.7.3 Covariant Field Equations

Since the cortex is not a *developable* surface, it cannot be given a parameterization whose metric tensor is uniform. We therefore developed an approach to take into account the intrinsic curvature of the solution domain when computing flow vector fields in the cortical parameter space. Similar approaches are common in fluid dynamics or general relativity applications, as they make sure that the mapping of one mesh surface onto another is parameterization-invariant. In the *covariant tensor* approach (Thompson et al., 2000), correction terms (Christoffel symbols, Γ^i_{jk} make the necessary adjustments for fluctuations in the metric tensor of the mapping procedure. In the partial differential equations (23), we replace L by the covariant differential operator $L^{\ddagger}$. In $L^{\ddagger}$, all L's partial derivatives are replaced with covariant derivatives. These covariant derivatives are defined with respect to the metric tensor of the surface domain where calculations are performed. The covariant derivative of a (contravariant) vector field, $u^i(\mathbf{x})$, is defined as $u^i_{,k} = \partial u^j/\partial x^k + \Gamma^j_{ik}$ where the *Christoffel symbols of the second kind* (Einstein, 1914), Γ^j_{ik}, are computed from derivatives of the metric tensor components $g_{jk}(\mathbf{x})$:

$$\Gamma^i_{jk} = (1/2)g^{il}(\partial g_{lj}/\partial x^k + \partial g_{lk}/\partial x^j + \partial g_{jk}/\partial x^i) \tag{1.24}$$

These correction terms are then used in the solution of the field equations to match one cortex with another. Note that a parameterization-invariant variational formulation could also be used to minimize metric distortion when mapping one surface to another. If P and Q are cortical surfaces with metric tensors $g_{jk}(u^i)$ and $h_{jk}(\xi^\alpha)$ in local coordinates u^i and ξ^α $(i, \alpha = 1, 2)$, the *Dirichlet energy* of the mapping $\xi(u)$ is defined as:

$$E(\xi) = \int_P e(\xi)(u)dP \tag{1.25}$$

where

$$e(\xi)(u) = g^{ij}(u)\partial\xi^{\alpha}(u)/\partial u^{i}\partial\xi^{\beta}(u)/\partial u^{j}h_{\alpha\beta}(\xi(u)) \tag{1.26}$$

and $dP = (\sqrt{\det[g_{ij}]}du^1du^2)$. The Euler equations, whose solution $\xi^{\alpha}(u)$ minimizes the mapping energy, are:

$$0 = L(\xi^i) = \sum_{m=1}^{2} \partial/\partial u^m[\det(g^{ru}) \sum_{l=1}^{2} g^{ml}_{ur}\partial\xi^i/\partial u^l] \tag{1.27}$$

(Liseikin, 1991). The resulting (harmonic) map (1) minimizes the change in metric from one surface to the other, and (2) is again independent of the parameterizations (spherical or planar) used for each surface. Intriguingly, similar variational approaches, using Beltrami flows, have been used for image restoration (Sochen et al., 1998), and related covariant PDEs can assist in structure extraction and statistical applications (see Thompson and Toga, 1998; Thompson et al., 2000 for a discussion).

1.8 Cortical Averaging

The warping fields deforming one cortex into gyral correspondence with another can also be used to create an *average* cortex (Fig. 14). First, all 38 gyral curves are transferred to the cortical parameter space, uniformly re-parameterized, and averaged to create an *average curve template* (*curves*, Fig. 13(b)). This serves as a target for alignment of individual cortical patterns (cf. Fischl et al., 1999). Each individual cortical pattern is transformed into the average curve configuration using a flow field in the parameter space (Fig. 13(a),(b)). By carrying a color code (that indexes 3D locations) along with the vector flow that aligns each individual with the average folding pattern, information can be recovered at a particular location in the average folding pattern (Fig. 13(d)) specifying the 3D cortical points mapping each subject to the average. This produces a new coordinate grid on a given subject's cortex (Fig. 13(f)) in which particular grid-points appear in the same location across subjects relative to the mean gyral pattern. By averaging these 3D positions across subjects, an average 3D cortical model is constructed for the group (Fig. 14, *bottom row*). The resulting mapping is guaranteed to average together all points falling on the same cortical locations across the set of brains, and ensures that corresponding features are averaged together.

1.8.1 Cortical Variability

By using the color code (Fig. 13(d)) to identify original cortical locations in 3D space (Fig. 13(f)), displacement fields are recovered mapping each patient into gyrus-by-gyrus correspondence with the average cortex (Fig. 15).

Anatomic variability is then defined at each point on the average cortex as the root mean square (r.m.s.) magnitude of the 3D displacement vectors, assigned to each point, in the surface maps driving individuals onto the group average (Thompson et al., 1996a,b, 1997, 1999). A typical variability pattern (based on 20 subjects) is visualized as a color-coded map in Fig. 16(a),(b). Overall, variability values rise sharply (Fig. 16(a),(b)) from 4-5 mm in primary motor cortex to localized peaks of maximal variability in posterior perisylvian zones, in the vicinity of visual area MT (or V5), and superior frontal association cortex (16-18 mm). Caution is therefore necessary when referring to functional activation foci or metabolic changes in this important area using stereotaxic coordinates, unless a non-linear registration approach is employed, otherwise structural differences may be interpreted as functional differences. The patterns also suggest a greater morphologic individuality in cortical regions that are phylogenetically more recent.

1.8.2 Average Brain Templates

Maps that deform individual cortical patterns into a group average shape can also assist in generating a brain template with the mean shape for a group. We recently used high-dimensional transformations to create a mean image template for a group of patients with Alzheimer's disease, whose anatomy is not well accommodated by existing brain atlases or imaging templates.

To make a mean image template for a group, several approaches are possible (Evans et al., 1994; Subsol, 1995; Grenander and Miller, 1998; Guimond et al., 1999; Thompson et al., 2000; Woods et al., 2000). If scans are mutually aligned using only a linear transformation (Fig. 17), the resulting average brain is blurred in the more variable anatomical regions, and cortical features are washed away. The resulting average brain also tends to exceed the average dimensions of the component brain images. By averaging geometric and intensity features separately (cf. Ge et al., 1995; Bookstein, 1997; Grenander and Miller, 1998; Christensen et al., 1999; Thompson et al., 2000), a template can be made with the mean intensity and geometry for a patient population. To illustrate this, we generated an initial image template for a group of Alzheimer's patients by (1) using automated linear transformations (Woods et al., 1993) to align the MRI data with a randomly selected image, (2) intensity-averaging the aligned scans, and then (3) recursively reregistering the scans to the resulting average affine image. The resulting average image was adjusted to have the mean affine shape for the group using matrix exponentiation to define average transformations (Woods et al., 1998). Images and a large set of anatomical surface models (84 per subject) were then linearly aligned to this template, and an average surface set was created for the group. Displacement maps driving the surface anatomy

of each subject into correspondence with the average surface set were then computed, and were extended to the full volume with surface-based elastic warping (Thompson et al., 2000). These warping fields reconfigured each subject's 3D image into the average anatomic configuration for the group. By averaging the reconfigured images (after intensity normalization), a crisp image template was created to represent the group (Fig. 17). Note the better-resolved cortical features and sharper definition of tissue boundaries in the average images after high-dimensional cortical registration. If desired, this AD-specific atlas can retain the coordinate matrix of the Talairach system (with the anterior commissure at (0,0,0)) while refining the gyral map of the Talairach atlas to encode the unique anatomy of the AD population. By explicitly computing matching fields that relate gyral patterns across subjects, a well-resolved and spatially consistent set of probabilistic anatomical models and average images can be made to represent the average anatomy and its variation in a subpopulation.

1.8.3 Uses of Average Templates

Average brain templates have a variety of uses. If functional imaging data from Alzheimer's patients is warped into an atlas template based on young normals, signals in regions with selective atrophy in disease are artificially expanded to match their scale in young normals, and biases can result. If the atlas has the average geometry for the diseased group, which may include atrophy, least distortion is applied by warping data into the atlas. Since the template (in Fig. 17) also has the average affine shape for the group (Woods et al., 1998), least distortion is applied when either linear, non-linear, approaches are used. The notion of least distortion can be formulated precisely using either (1) the associated matrix and deformation tensor metrics (Woods et al., 2000), or (2) using the L^2-norm on the Hilbert space of deformation field coefficients (Grenander and Miller, 1998), or (3) indirectly through a continuum-mechanical operator or regularization functional that defines what it means for a distortion to be irregular (Christensen et al., 1999; see Section 17.4). Interestingly, automated registration approaches were able to reduce anatomic variability to a greater degree if a specially-prepared image template was used as a registration target (Thompson et al., 1999; Woods et al., 1999, 2000). With smaller deformations, non-global minima of the registration measure may be avoided, and convergence may also be faster, as the parameter space is searched for an optimal match. Average templates are under rapid development for the *Macaque* brain (Grenander and Miller, 1998), and for individual structures such as the *corpus callosum*, (Davatzikos, 1996; Gee et al., 1998), central sulcus (Manceaux-Demiau et al., 1998), cingulate and paracingulate sulci (Paus et al., 1996; Thompson et al., 1997), hippocampus (Haller et al., 1997; Joshi et al., 1998; Csernansky

et al., 1998; Thompson et al., 1999) and for transformed representations of the human and *Macaque* cortex (Drury and Van Essen, 1997; Grenander and Miller, 1998; Thompson et al., 1999; Fischl et al., 1999).

1.9 Deformation-Based Morphometry

1.9.1 Deformable Probabilistic Atlases

As noted earlier, *warping* algorithms create deformation maps (Fig. 8) that indicate 3D patterns of anatomic differences between any pair of subjects. By defining probability distributions on the space of deformation transformations which drive the anatomy of different subjects into correspondence (Grenander, 1976; Amit et al., 1991; Grenander and Miller, 1994; Thompson and Toga, 1997; Thompson et al., 1997), statistical parameters of these distributions can be estimated from databased anatomic data to determine the magnitude and directional biases of anatomic variation. Encoding of local variation can then be used to assess the severity of structural variants outside of the normal range, which, in brain data, may be a sign of disease (Thompson et al., 1997). Methods for representing variation and detecting group difference in shape can be collectively referred to as approaches for *deformation-based morphometry.*

1.9.2 Encoding Brain Variation

To see if disease-specific features could be detected in individual patients, we developed a random vector field approach to construct a population-based brain atlas (Thompson and Toga, 1997). Briefly, given a 3D MR image of a new subject, a warping algorithm calculates a set of high-dimensional volumetric maps, elastically matching this image with other scans from an anatomic image database. Target scans are selected from subjects matched for age, handedness, gender, and other demographic factors (Thompson et al., 1997, 1998). The resulting family of volumetric warps provides empirical information on local variability patterns. A probability space of random transformations, based on the theory of anisotropic Gaussian random fields (Thompson et al., 1997), is then used to encode the variations. As noted earlier (Section 17.7), cortical matching approaches are needed to represent variations in gyral patterns (Thompson et al., 1997; Thompson and Toga, 1998). Confidence limits in stereotaxic space are determined, for points in the new subject's brain, enabling the creation of color-coded probability maps to highlight and quantify regional patterns of deformity (Fig. 18).

1.9.3 Tensor Maps of Directional Variation

Fig. 16 shows a tensor map of variability for normal subjects, after mapping 20 elderly subjects' data into Talairach space (all right handed, 10 males, 10 females). Rectangular glyphs indicate the principal directions of

variation - they are most elongated along directions where anatomic variation is greatest across subjects. Clearly, structures do not vary to the same degree in every coordinate direction (Thompson et al., 1996), and even these directional biases vary by cortical system. Each glyph represents the covariance tensor of the vector fields that map individual subjects onto their group average. Because gyral patterns constrain the mappings, the fields reflect variations in cortical organization at a more local level than can be achieved by only matching global cortical geometry. Note the elongated glyphs in anterior temporal cortex, and the very low variability (in any direction) in entorhinal and inferior frontal areas. By better defining the parameters of allowable normal variations, the resulting information can be leveraged to distinguish normal from abnormal anatomical variants. The tensor fields make it easier to detect anomalies, which may be small in magnitude but in an unusual direction. Local computation of the variance components also means that confidence limits for abnormal structure are appropriately relaxed in regions of high anatomic variability, so normal differences are not signaled as deficits.

1.9.4 Anisotropic Gaussian Fields

In a probabilistic atlas, well-defined statistical criteria are required to identify significant differences in brain structure. These criteria can be formulated in different ways, depending on the attribute whose statistical variation is being modeled. One approach is to use the theory of Gaussian random fields, a modeling technique used widely in functional image analysis (e.g., SPM; Friston et al., 1995). By contrast with functional signals, which are generally treated as random *scalar* fields, the deformation maps that quantify structural differences are treated as random *vector* fields. Instead of a field of variance values, the variability of the deformation vectors, and their directional tendencies, are stored using a covariance tensor at each anatomical point (Thompson et al., 1996; Cao and Worsley, 1999).

In one study (Thompson et al., 1997; cf. Cao and Worsley, 1999), we developed an approach to detect brain structure differences between two groups, or between an individual subject and a database of demographically matched subjects. Suppose $\mathbf{W}_{ij}(\mathbf{x})$ is the deformation vector required to match the structure at position $\mathbf{x}$ in an atlas template with its counterpart in subject i of group j. (If surface models are being analyzed, rather than full brain volumes (Thompson and Toga, 1998), $\mathbf{W}_{ij}(\mathbf{x})$ is the deformation vector matching parametric mesh node $\mathbf{x}(u, v)$ with its counterpart in subject i of group j.) We then model the deformations as:

$$\mathbf{W}_{ij}(\mathbf{x}) = \mu_{\mathbf{j}}(\mathbf{x}) + \mathbf{\Sigma}(\mathbf{x})^{\mathbf{1/2}}\epsilon_{\mathbf{ij}}(\mathbf{x}) \tag{1.28}$$

Here $\mu_{\mathbf{j}}(\mathbf{x})$ is the mean deformation for group j, and $\mathbf{\Sigma}(\mathbf{x})$ is a non-stationary,

anisotropic covariance tensor field, which relaxes the confidence threshold for detecting abnormal structure in regions where normal variability is extreme, $\mathbf{\Sigma}(\mathbf{x})^{1/2}$ is the upper triangular Cholesky factor tensor field (Fig. 16), and $\epsilon_{\mathbf{ij}}(\mathbf{x})$ is a trivariate random vector field whose components are independent stationary Gaussian random fields.

1.9.5 Detecting Shape Differences

A T^2 or F statistic that indicates evidence of significant difference in deformations between the groups is calculated at each lattice location in a 3D image or parameterized 3D surface, to form a statistic image. Under the null hypothesis of no abnormal deformations, the statistic image is approximated by a T^2 random field. Specifically, the significance of a difference in brain structure between two subject groups (e.g., patients and controls) of N_1 and N_2 subjects is assessed by calculating the sample mean and variance of the deformation fields ($j = 1, 2$):

$$\mathbf{W}^{\mu}_{\mathbf{j}}(\mathbf{x}) = \sum_{i=1}^{N_j} \mathbf{W_{ij}}(\mathbf{x})/\mathbf{N_j} \quad (1.29)$$

$$\mathbf{\Psi}(\mathbf{x}) = \frac{1}{N_1 + N_2 + 2}\{\sum_{j=1}^{2}\sum_{i=1}^{N_j}[\mathbf{W_{ij}}(\mathbf{x}) - \mathbf{W}^{\mu}_{\mathbf{j}}(\mathbf{x})][\mathbf{W_{ij}}(\mathbf{x}) - \mathbf{W}^{\mu}_{\mathbf{j}}(\mathbf{x})]^{\mathbf{T}}\} \quad (1.30)$$

and computing the following statistical map (Thompson et al, 1997; Cao and Worsley, 1999):

$$T^2(\mathbf{x}) = \{\frac{N_1 N_2}{(N_1 + N_2)(N_1 + N_2 - 2)}\}[\mathbf{W}^{\mu}_{\mathbf{2}}(\mathbf{x}) - \mathbf{W}^{\mu}_{\mathbf{1}}(\mathbf{x})]^{\mathbf{T}}\mathbf{\Psi}^{-\mathbf{1}}(\mathbf{x})[\mathbf{W}^{\mu}_{\mathbf{2}}(\mathbf{x}) - \mathbf{W}^{\mu}_{\mathbf{1}}(\mathbf{x})] \quad (1.31)$$

Under the null hypothesis, $(N_1 + N_2 - 2)T^2(\mathbf{x})$ is a stationary Hotelling's T^2-distributed random field. At each point, if we let $\nu = (N_1 + N_2 - 2)$ and we let the dimension of the search space be d=3, then:

$$F(\mathbf{x}) = ((\nu - d + 1)/d)T^2(\mathbf{x}) \sim F_{d,(\nu-d+1)} \quad (1.32)$$

In other words, the field can be transformed point-wise to a Fisher-Snedecor F distribution (Thompson et al., 1997). To obtain a p-value for the effect that is adjusted for the multiple comparisons involved in assessing a whole field of statistics, Cao and Worsley (1999) examined the distribution of the global maximum T^2_{max} of the resulting T^2-distributed random field under the null hypothesis. The resulting probability that $T^2(\mathbf{x})$ ever exceeds a fixed high threshold T^2_{max} is approximated by the expected Euler

characteristic $E[\chi(A(T^2_{max}))]$ of the excursion sets of the Hotelling's T^2-distributed random field above the threshold T^2_{max}. Then $p[T^2_{max} \geq t]$ is approximated by $\sum_{n=0}^{d} R_n \rho_n(t)$, where the number of n-dimensional resolution elements $R_n = V_n/(FWHM)^n$ depends on the effective full-width-at-half-max (FWHM) of the component Gaussian images $\epsilon_{\mathbf{ij}}(\mathbf{x})$, and on the Euler characteristic (V_0), caliper diameter ($V_1/2$), surface area ($2V_2$) and volume (V_3) of the search region. The n-dimensional EC densities are given by (Cao and Worsley, 1999):

$$\rho_0(t) = \int_t^\infty [\frac{\Gamma((\nu+1)/2)}{(\nu\pi)^{1/2}\Gamma(\nu/2)}].[1+(u^2/\nu)]^{-1/2(\nu+1)} du \tag{1.33}$$

$$\rho_1(t) = \frac{(4\ln 2)^{1/2}}{2\pi}.[1+(t^2/\nu)]^{-1/2(\nu-1)} \tag{1.34}$$

$$\rho_2(t) = \frac{(4\ln 2)}{(2\pi)^{3/2}}.[\frac{\Gamma((\nu+1)/2)}{\nu/2^{1/2}\Gamma((\nu)/2)}].t[1+(t^2/\nu)]^{-1/2(\nu-1)} \tag{1.35}$$

$$\rho_3(t) = \frac{(4\ln 2)^{3/2}}{(2\pi)^2}.[((\nu-1)/\nu)t^2-1].[1+(t^2/\nu)]^{-1/2(\nu-1)} \tag{1.36}$$

The global maximum of the random deformation field, or derived tensor fields (Thompson et al., 1998; *see below*), can be used to test the hypothesis of no structural change in disease (Worsley, 1994a,b; Cao and Worsley, 1999). Similar multivariate linear models can be used to test for the effect of explanatory variables (e.g., age, gender, clinical test scores) on a set of deformation field images (Ashburner et al., 1998; Gaser et al., 1998). This can help explore linkages between atlas descriptions of variance and behavioral or cognitive parameters (Fuh et al., 1997; Mega et al., 1998; Zoumalan et al., 1999)

1.9.6 Tensor-Based Morphometry

Since the goal of brain mapping is to localize structural and functional effects in the brain, local changes must be distinguished from positional shifts of a given brain region. This can be a confounding problem when structural or functional data are analyzed voxel-by-voxel after only a linear transformation into stereotaxic space, since the tissue that appears in a given stereotaxic voxel may shift. This may be due to a structural difference in a distant anatomical region, which moves the relative positions of other tissues, or it may be due to a registration difference where a distant anatomical change alters the ability of the algorithm to align other image voxels. Ideally then, disease-related shape changes that affect a given structure in the brain should be separated into dilations and contractions within the structure itself, and global mechanical or positional shifts due to changes

in other brain structures. Tensor maps (Thompson et al., 2000) offer a solution to this problem. Because they refer to the local derivatives of the deformation field, tensor maps of shape differences are invariant to global translational differences in structures. The shape differences they detect are therefore less sensitive to registration artifacts. To characterize purely local differences, tensor maps examine differential properties of the deformation fields, including local expansion, contraction or shearing effects recovered by the warping transformation. The Jacobian of the deformation field has been used as a local index of gender-specific shape differences in the corpus callosum (Davatzikos, 1996; Machado and Gee, 1998). Other local vector field operators, including the gradient and divergence ($\nabla \mathbf{u}(\mathbf{x}), \nabla^T \bullet \mathbf{u}(\mathbf{x})$; Subsol et al., 1997; Thompson et al., 1998), the specialized norm $\times$ divergence operator ($\| \mathbf{u}(\mathbf{x}) \| \nabla^T \bullet \mathbf{u}(\mathbf{x})$; Thirion and Calmon, 1997; Thirion et al., 1998) and the delta-filter (Bookstein, 1997) have been applied to deformation fields, in studies of brain development, tumor progression, and multiple sclerosis lesion growth. Each of these operators is designed to emphasize different aspects of growth or pathologic processes. Their differing signal-to-noise therefore offers an approach to enhance the detection sensitivity of deformation-based morphometry.

1.9.7 Mapping Brain Asymmetry

There is a vast literature on cortical surface asymmetries (Eberstaller, 1884; Cunningham, 1892; Geschwind and Levitsky, 1968; Davidson and Hugdahl, 1994). These asymmetries have been related to functional lateralization (Strauss et al., 1983), handedness (Witelson, 1989), language function (Davidson and Hugdahl, 1994), and asymmetries of associated cytoarchitectonic fields (Galaburda and Geschwind, 1981) and their thalamic projection areas (Eidelberg and Galaburda, 1982). After group averaging of anatomy, asymmetric features emerge that are not observed in individual anatomies due to their considerable variability. As shown in Fig. 16 (*sagittal projection*), the marked anatomic asymmetry in posterior perisylvian cortex, actually extends rostrally into postcentral cortex, with the posterior bank of the postcentral gyrus thrust forward by 8-9 mm on the right compared to the left. The asymmetry also extends caudally across the lateral convexity into superior and inferior temporal cortex. As shown earlier by averaging ventricular models (Fig. 5), this asymmetry penetrates subcortically into the occipital ventricular horn, but not into adjacent parieto-occipital and calcarine cortex (Thompson et al., 1997).

The improved ability to localize asymmetry and encode its variability in a disease-specific atlas has interesting applications in schizophrenia (Narr et al., 1999). Schizophrenic patients have anatomic alterations in several brain regions, including the superior temporal gyrus (e.g., Nestor et al.,

1993) and several studies have reported a lack of asymmetry in schizophrenic patients (e.g., Kikinis et al., 1994) although the findings remain controversial (Frangou et al., 1997).

1.9.8 Changes in Asymmetry

To see if cortical asymmetries were lost in schizophrenia, we made average cortical representations for schizophrenic patients (15 males, all right-handed) and matched controls (also 15 males, right-handed; Narr et al., 2000). As described in Section 17.8, 38 major sulcal curves were used to drive each subject's gyral pattern into a group mean configuration (Fig. 19). The magnitude of anatomic variation in each brain region was also computed from the deformation vector fields, and shown in color as a variability map (Fig. 19, *colors*). Perhaps surprisingly, asymmetry was not attenuated in the patient group. This can be seen immediately in the sagittal projections of average anatomy for each group. Significant asymmetries were confirmed by calculating curvature and extent measures from the parametric mesh models (Narr et al., 2000). In frontal cortex, the patients also displayed greater variability than controls. Since relatively subtle asymmetries emerge clearly in a group atlas, these disease-specific, population-based atlases may be advantageous for investigating a variety of alterations in cortical organization or lateralization, and their dependencies on genetic parameters (Thompson et al., 2000a,b; cf. Csernansky et al., 1998; Le Goualher et al., 1999). In a related study (Sowell et al., 2000), we found a robust and significant increase in the average magnitude of cortical asymmetry across the first 30 years of life. There was also a further increase in average cortical asymmetry between normal elderly subjects and Alzheimer's patients (Thompson et al., 1998), in regions where asymmetric changes in perfusion and metabolism occur early in the disease.

1.9.9 Abnormal Asymmetry

In an interesting development, Thirion et al. (1998) applied a warping algorithm to a range of subjects' scans, in each case matching each brain hemisphere with a reflected version of the opposite hemisphere. The resulting asymmetry fields were treated as observations from a spatially-parameterized random vector field, and deviations due to lesion growth or ventricular enlargement were detected using the theory developed in (Thompson et al., 1997). Due to the asymmetric progression of many degenerative disorders (Thompson et al., 1998), abnormal asymmetry may prove to be a sensitive index of pathology in individual subjects or groups. From a more practical standpoint, asymmetry fields are smaller in magnitude than subject-to-subject deformation maps. This makes the fields easier to estimate with automated non-linear registration algorithms. When the estimated deforma-

tion is small, it is easier to avoid false, non-global minima of the matching measure being optimized. This improves the robustness of deformation-based morphometry for pathology detection.

1.9.10 Model-Based Shape Analysis

Mesh-based modeling of anatomy (Thompson et al., 1996a,b, 1997; Narr et al., 2000) allows detailed analysis of morphometric differences for specific structures. In recent years, we have developed an anatomical modeling approach which imposes a parametric computational grid onto anatomical surfaces (up to 84 per brain). Shape parameters are then computed, including surface complexity, curvature and asymmetry indices, as well as surface area and extent measures in 3 dimensions (Thompson et al., 1996). These shape measures are then typically subjected to multivariate analysis of variance (MANOVA), to test hypotheses about disease-specific changes. Interactions between disease and gender, age and other demographic factors can also be assessed (Thompson et al., 2000). Significantly altered structural patterns were detected in patients with, or at risk for, Alzheimer's Disease (Thompson et al., 1997; Kwong et al., 1999), in adult and childhood-onset schizophrenia (Narr et al., 2000; Blanton et al., 1999), and in fetal alcohol syndrome (Sowell et al., 2000). In dementia, significant linear relationships were also found between localized shape differences and cognitive scores (Mega et al., 1998). The mesh-based approach has also uncovered gender differences in the development subcortical structures (Blanton et al., 1999), interactions between gender and disease (Narr et al., 2000) and subtle developmental shape changes in childhood (Thompson et al., 2000; Blanton et al., 2000), and adolescence (Sowell et al., 2000). The resulting toolkit (Thompson et al., 1996, 1998, 2000) provides an approach for the statistical analysis of shape, and creates average models, maps of variability, asymmetry, and group differences that can be used to build disease-specific brain atlases (Thompson et al., 2000a,b). Related approaches for neuroanatomical shape analysis, based on Riemannian shape manifolds (Bookstein, 1997) and pattern theory (Grenander and Miller, 1998), are reviewed in (Thompson and Toga, 1998; Thompson et al., 2000).

1.10 Voxel-Based Morphometry

1.10.1 Detecting Changes in Stereotaxic Tissue Distribution

Additional techniques have been applied to detect group differences in the stereotaxic distribution of tissue types in the brain. In these approaches, binary maps of gray matter, white matter and CSF are analyzed, after linear alignment of individual MRI scans to an image template in Talairach space (Andreasen et al., 1994; Wright et al., 1994; Sowell et al., 1999a,b; Paus et al., 1999). Aligned volumes are first RF-corrected and segmented, using a

tissue classifier, to produce binary maps for each tissue type. Differences in incidence of tissue, at each stereotaxic voxel, are assessed by computing a statistical parametric map (SPM; Friston et al., 1995) in which each voxel contains a statistic quantifying the group difference at that stereotaxic position. Because this approach investigates group differences in tissue incidence at each stereotaxic voxel, it is commonly referred to as *voxel-based morphometry*. At each voxel, the actual statistic is compared with a reference (or *null*) distribution for the statistic (i.e. the values it takes when groups are sampled from the same population, and no effect is present). This gives a p-value for how likely it is that such a difference could occur by accident, if only that voxel were assessed. Because a vast number of voxels are assessed, p-values have to be corrected for multiple comparisons before the significance of the results are assessed, unless there was an *a priori* hypothesis of an effect at a specific stereotaxic voxel. Bonferroni corrections, which adjust p-values based on the total number of independent tests, are not used because data at neighboring voxels are highly correlated. Approaches to obtain corrected p-values include the theory of stationary Gaussian random fields (Wright et al., 1994), statistical flattening (Worsley et al., 1999), and permutation (Sowell et al., 1999a,b). These approaches are described next.

1.10.2 Stationary Gaussian Random Fields

The application of Gaussian field theory to detect significant group differences in brain images has revolutionized the field of brain mapping (Friston et al., 1995; Frackowiak et al., 1997). Typically applied to functional images of brain activation or metabolism, Gaussian field theory models the distributions of features in statistical maps that would be found by accident, if the null hypothesis of no significant difference between groups were true. Experimental effects are compared with these null distributions, to check if they could have occurred accidentally, or whether there is enough evidence to reject the null hypothesis of no differences between groups. The features whose distributions are modeled include (1) the maximum value (or peak height, Z_{max}) of the statistic that would be found in the map, and (2) the size of the largest connected cluster of voxels above a given threshold. Null distributions for more complex features can also be derived mathematically, such as the number of clusters exceeding a given height and spatial extent, or the total spatial extent of these clusters. These features are thought of as measurements of 'rising swells or waves in a choppy (noisy) sea' (Lange, 1998), where the roughness of the sea is estimated from the data. To estimate the probability that the maximum value of the map (Z_{max}) is greater than a given threshold t under the null hypothesis (i.e., when no difference is present), Worsley et al. (1994) used the expected Euler characteristic $E[\chi(A(t))]$ of a binarized map thresholded at t, so that for high t,

$$Pr(Z_{max} \succ t) \cong E[\chi(A(t))] = \lambda(V)|\Lambda|^{1/2}(2\pi)^{(D+1)/2}.He_D(t).exp(-t^2/2) \tag{1.37}$$

Here $\lambda(V)$ and D are the volume and dimension of the search region, and $He_D(t)$ is the D-th order Hermite polynomial. The roughness tensor, Λ (or its inverse, the smoothness tensor, Λ^{-1}), is crucial for estimating p-values. It is defined as the covariance matrix of the partial derivatives of the process along each of the D coordinate axes, with variances $Var[\partial X/\partial x_i]$ on the diagonal and off-diagonal elements $Cov[\partial X/\partial x_i, \partial X/\partial x_j]$. Usually the smoothness is calculated not from the data itself, which may contain a physiological signal, but from the residuals after fitting a linear statistical model which removes linear effects of the experimental parameters. For white noise smoothed by a Gaussian filter with full-width at half-max $FWHM_i$ along the i-th axis, the smoothness is related to the filter parameters:

$$|\Lambda|^{-1/2} = (4\ln 2)^{-D/2}.\prod_{i=1}^{D} FWHM_i \tag{1.38}$$

Once these parameters are estimated, a significance level can be assigned to the overall experiment, so long as the theoretical assumptions are not violated.

1.10.3 Statistical Flattening

An underlying assumption of the above parametric approach is that the process is a *stationary* Gaussian field, i.e. its statistical characteristics, including its roughness parameter $|\Lambda|$ (or its reciprocal, the smoothness $|\Lambda|^{-1}$), are the same at each point in the image. The FWHM of the process should be constant in all directions, and across all voxels in the image. While these assumptions are reasonable for functional imaging data, they are likely to be violated for structural imaging data. Binary structure masks, for example, are constant across large regions, and even after smoothing the signal changes more rapidly at the edges of structures (Worsley et al., 1999). The distribution of cluster sizes that occur by accident is therefore considerably skewed towards larger cluster sizes in smooth regions of the image, resulting in more false positives (and false negatives in rough regions) than predicted by formulae for stationary fields. To address this, Worsley et al., (1999) suggested a *statistical flattening* approach in which the data are warped into a new space, which may have higher dimension than the data, so that in the new space the smoothness of the normalized residuals of the statistical model is stationary. The p-value for cluster sizes above a threshold can then be applied using size measurements in the new space, or by estimating the effective resolution of the field directly from the normalized residuals (Worsley et al., 1999).

1.10.4 Permutation

A final approach to estimating p-values for significant features in statistical maps is to estimate their distribution under the null hypothesis by permutation (Sowell et al., 1999a,b). This non-parametric approach avoids assumptions about the spatial autocorrelation of the process, and has been successful in functional imaging as well (Holmes et al., 1996). Subjects are randomly assigned to groups and the distribution of accidental clusters is tabulated empirically. In a recent study of maturation changes in gray matter distribution between childhood and adolescence (Sowell et al., 1999), we found significant regional reductions in gray matter that were specific to dorsal frontal and parietal cortices (Fig. 20; $p \prec 0.05$, *permutation test*). Although random permutations revealed that false positive clusters occurred (on average 5.8 per simulation), the age-contrast revealed a number of suprathreshold clusters (57) which was significantly higher than predicted by the null distribution.

1.10.5 Joint Assessment of Shape and Tissue Distribution

Non-linear registration provides an opportunity to integrate the approaches for deformation-based and voxel-based morphometry (Thompson et al., 2000). By jointly assessing group differences in anatomical shape and tissue distribution, powerful tests can be developed to detect disease-specific differences. In dementia for example, it is of practical value to track the profile of cortical gray matter loss in an individual subject in vivo. Nonetheless, the extreme variability in gyral patterns confounds efforts to calibrate this loss against a normative population, or determine the average profile of early tissue loss in a group. With an approach to compute gyral pattern variations across subjects, variations in cortical gray matter distribution can be mathematically separated from variations in cortical organization, and both can be visualized and statistically assessed. Fig. 22 shows the profile of cortical gray matter loss in a population of patients with mild to moderate Alzheimer's disease, relative to a group of elderly conrol subjects matched for age, gender, handedness and educational level. Measurements of gray matter density were derived from filtered gray matter maps, and textured onto cortical models as an attribute for each subject. By using cortical pattern matching to define corresponding regions across subjects, gray matter measurements could be averaged and compared across homologous regions of cortex for all 46 subjects in the study. When these patterns of cortical variability were directly encoded using a random vector field, and a secondary statistical distribution was used to encode gray matter variation in normals, severe reductions in gray matter (up to 30% loss, greater in the left hemisphere) were observed across the lateral temporal surfaces in the AD cohort. This pattern is consistent with the profile of early perfusion deficits

and metabolic reductions. If an *a priori* hypothesis had not been available on the localization of gray matter reductions, corrected p-values for the reduction could be obtained by permutation. A variational approach can also be developed to assess the statistical significance of gray matter loss across the cortical sheet. If a partial differential equation:

$$g^{ij}(\partial^2 \mathbf{u}/\partial r^i \partial r^j) + \partial/\partial u^j (S^{ij})\mathbf{u_{r^i}} = 0 \tag{1.39}$$

is run in the parameter space of the group average cortex, this generates a deformed grid $\mathbf{u}(\mathbf{r})$ whose deformation gradient tensor approximates the smoothness tensor S^{ij} of the normalized residuals of the gray matter distribution on the surface (here g^{ij} is the contravariant metric tensor of the grid). If the smoothness tensor has non-zero curvature (and is therefore not realizable as a deformation tensor), the deformation gradient approximates it in the Frobenius matrix norm. Relative to this new computational grid, the residuals become stationary and isotropic, and p-values for the gray matter reductions can be evaluated.

1.11 Dynamic (4D) Brain Maps

If the brain is imaged repeatedly over a period of time, dynamic changes in brain structure can be investigated. These structural changes include tumor growth, degenerative change, and the growth patterns of normal and abnormal brain development. Static representations of brain structure may not be ideal for determining the complex dynamics of brain development and disease. In developmental disorders, for example, a child may display a normal phenotype with an aberrant time-course. Similarly, if growth rates are abnormal, morphology may not be detectably different at any time-point, due to the wide variations in normal anatomy. Ideally, the dynamics of brain growth and degeneration should be compared against a database of dynamic normative data.

Using high-dimensional surface-based warping, we recently identified an anterior-to-posterior wave of peak growth rates at the *corpus callosum*, during the first 15 years of life (Fig. 22; Thompson et al., 2000). Scans were acquired from young normal children across multi-year time-spans (up to 4 years), and maps of local growth rates were derived from elastic transformations that reconfigured the earlier anatomy into its later configuration. By applying local operators to the deformation fields, *tensor* maps were created to reflect the magnitude and principal directions of tissue dilation or contraction, and the local rates, divergence and gradients of the growth processes detected in the dynamically changing brain. The reliability of the results was also investigated by acquiring null maps of anatomical change across very short intervals. After the age of 6, peak growth rates were consistently found in regions of the *corpus callosum* that connect linguistic and associa-

tion cortices of the two brain hemispheres. After puberty, these growth rates were considerably reduced, and tissue loss was also identified in subcortical regions.

Growth maps offer improved spatial detail and detection sensitivity that may be advantageous in detecting early degenerative change in dementia. They may also be used to track the effects of therapeutic interventions in patients with tumor growth, active lesions, and traumatic brain injury. The ability to map the local dynamics of growth in an individual child is also advantageous scientifically and clinically. As a result, we are currently applying this approach to wider normative subject and patient populations to establish dynamic criteria for assessing developmental and degenerative disease processes.

1.12 Conclusion

In this chapter, we reviewed some exciting developments in the field of brain image analysis and atlas construction. Brain imaging studies are expanding into ever-larger populations, and this enables digital atlases to be developed that synthesize brain data across vast numbers of subjects. Data can also be fused across multiple imaging modalities, to better understand brain structure and function in health and disease. Mathematical algorithms can exploit the data in these population-based atlases to detect pathology in an individual or patient group, to detect group features of anatomy not apparent in an individual, and to uncover powerful linkages between structure and demographic or genetic parameters. These algorithms in turn draw upon developments in the computer vision, pattern recognition, graphics, and statistical research arenas. Above all, the pace of technological development in brain imaging has clearly been matched by the rapid emergence of powerful algorithms for data analysis. The resulting armory of tools shows enormous promise in shedding light on the complex structural and functional organization of the human brain.

1.13 Acknowledgments

This work was supported by research grants from the National Center for Research Resources (P41 RR13642 and RR05956), the National Institute of Neurological Disorders and Stroke and the National Institute of Mental Health (NINDS/NIMH NS38753), and by a Human Brain Project grant to the International Consortium for Brain Mapping, funded jointly by NIMH and NIDA (P20 MH/DA52176). Additional support was also provided by the National Library of Medicine (LM/MH05639) and National Science Foundation (BIR 93-22434). Special thanks go to our colleagues Roger Woods, Jay Giedd, Colin Holmes, David MacDonald, Alan Evans, and John Mazziotta, whose support has been invaluable in these investigations.

.....