

A probabilistic atlas and reference system for the human brain: International Consortium for Brain Mapping (ICBM)

John Mazziotta^{1*}, Arthur Toga², Alan Evans³, Peter Fox⁴, Jack Lancaster⁴, Karl Zilles^{5,11}, Roger Woods¹, Tomas Paus³, Gregory Simpson⁶, Bruce Pike³, Colin Holmes⁸, Louis Collins³, Paul Thompson², David MacDonald³, Marco Iacoboni¹, Thorsten Schormann¹³, Katrin Amunts¹⁴, Nicola Palomero-Gallagher¹², Stefan Geyer¹⁵, Larry Parsons⁴, Katherine Narr², Noor Kabani³, Georges Le Goualher³, Dorret Boomsma⁹, Tyrone Cannon⁷, Ryuta Kawashima¹¹ and Bernard Mazoyer¹⁰

¹Ahmanson-Lovelace Brain Mapping Center, UCLA School of Medicine, 660 Charles E. Young Drive, South Los Angeles, CA 90095, USA ²Laboratory of Neuroimaging, Department of Neurology, UCLA School of Medicine, 710 Westwood Blvd, Los Angeles, CA 90095, USA ³McGill University/McConnell Brain Imaging Center, Montreal Neurological Institute, WB-315, 3801 University Street, Montreal, Quebec H3A 2B4, Canada ⁴Research Imaging Center, University of Texas Health Sciences Center at San Antonio, 7703 Floyd Curl Drive, San Antonio, TX 78284, USA ⁵University Duesseldorf, Vogt Brain Research Institute, P.O. Box 10 10 07, Duesseldorf D-40001, Germany ⁶Dynamic NeuroImaging Laboratory, UCSF, Department of Radiology, Box 0628, San Francisco, CA 94143, USA ⁷UCLA Department of Psychology, Franz Hall, 405 Hilgard Avenue, Los Angeles, CA 90095-1563 ⁸SGL, 43L-525, 1600 Amphitheatre Pkwy, Mountain View, CA 94043-1352, USA ⁹Vrije Universiteit, Department of Biological Psychology, De Boelelaan 1111, 1081 HV Amsterdam ¹⁰Groupe d'Imagerie Neurofonctionnelle, UMR 6095 CNRS, CEA, Université de Caen, Université Paris 5, GIP Cyceron, BP5229, F-14074 Caen, France ¹¹Department of Nuclear Medicine and Radiology, Institute of Development, Aging and Cancer, Tohoku University, 4-1 Seiryomachi Aobaku, Sendai 980-8575, Japan ¹²Institute of Medicine, Research Centre, Jülich, D-52425, Germany ¹³Institut für Neuroanatomie und C. und O., Vogt Institut für Hirnforschung, Heinrich-Heine-Universität, Postfach 101007, D-40001 Düsseldorf, Germany ¹⁴Institut für Medizin, Forschungszentrum Jülich und Institut für Neuroanatomie, Heinrich-Heine-Universität, Postfach 101007, D-40001 Düsseldorf, Germany ¹⁵Institute of Neuroanatomy, Heinrich Heine University, Moorenstr. 5, D-40225 Düsseldorf, Germany

Motivated by the vast amount of information that is rapidly accumulating about the human brain in digital form, we embarked upon a program in 1992 to develop a four-dimensional probabilistic atlas and reference system for the human brain. Through an International Consortium for Brain Mapping (ICBM) a dataset is being collected that includes 7000 subjects between the ages of eighteen and ninety years and including 342 mono- and dizygotic twins. Data on each subject includes detailed demographic, clinical, behavioural and imaging information. DNA has been collected for genotyping from 5800 subjects. A component of the programme uses post-mortem tissue to determine the probabilistic distribution of microscopic cyto- and chemoarchitectural regions in the human brain. This, combined with macroscopic information about structure and function derived from subjects *in vivo*, provides the first large scale opportunity to gain meaningful insights into the concordance or discordance in micro- and macroscopic structure and function. The philosophy, strategy, algorithm development, data acquisition techniques and validation methods are described in this report along with database structures. Examples of results are described for the normal adult human brain as well as examples in patients with Alzheimer's disease and multiple sclerosis. The ability to quantify the variance of the human brain as a function of age in a large population of subjects for whom data is also available about their genetic composition and behaviour will allow for the first assessment of cerebral genotype-phenotype-behavioural correlations in humans to take place in a population this large. This approach and its application should provide new insights and opportunities for investigators interested in basic neuroscience, clinical diagnostics and the evaluation of neuropsychiatric disorders in patients.

Keywords: atlas; probabilistic; cytoarchitecture; magnetic resonance imaging; neuroanatomy; genetics

1. INTRODUCTION

The nervous system is unique among human body systems in its spatial and temporal organization. The central nervous system is divided into highly specialized

regions that have unique properties in terms of cell types, connections and organization. The functions of these units vary with time, spanning the gamut from the millennia of evolution to the millisecond choreography of neurophysiological events. This temporal and spatial specialization is well suited to the application of informatics techniques. In fact, such methods will be required

*Author for correspondence (mazz@loni.ucla.edu).

as the basis from which to begin to understand and organize the ever increasing amount of neuroscientific information that is accumulating about this, the most complicated system known. For each brain region and for every attribute ascribed to it, data must be organized in a rational fashion that takes advantage of the brain's inherent neuroanatomical organization, spatial segregation and the variance among individuals that exist for these regions. What is ultimately required is a multidimensional database organized with three dimensions in space and one in time, along with a seemingly infinite number of attributes referable to these four physical dimensions.

As in geography, neuroscience requires accepted maps, terminologies, coordinate systems and reference spaces to allow accurate and effective communication within the field and to allied disciplines. Geographical atlases of the earth have advantages over atlases that are anatomical in nature. Earth atlases can assume a relatively constant physical representation over thousands of years. On that single, stable construct, an infinite number of abstract representations of features can be overlaid. For earth maps, such features might include rainfall, temperature, population density or crime rates.

Unlike geographical atlases, anatomical atlases cannot assume a single, constant physical reality. This is true despite the fact that standard atlases that utilized single subjects minimize this fundamental problem. Anatomical atlases must first deal with the fact that there is a potentially infinite number of physical realities that must be modelled to obtain an accurate, probabilistic representation of the entire population. Upon this anatomical representation one can then overlay features in a fashion analogous to that described for earth atlases (Mazziotta 1984). In the brain, such features might include cytoarchitecture, chemoarchitecture, blood flow distributions, metabolic rates, ligand binding, behavioural and pathological correlates, and many others. Like earth maps, brain maps can vary in time frames ranging from milliseconds (e.g. electrophysiological events) to minutes (e.g. skill acquisition), years (e.g. development, maturation, ageing), or millennia (i.e. evolution).

Classical atlases of the human brain or other species have been derived from a single brain, or brains from a very small number of subjects, and have employed simple scaling factors to stretch or constrict a given subject's brain to match the atlas. The result is a rigid and often inflexible system that disregards useful information about morphometric (i.e. dimensionality) and densitometric (i.e. intensity) variability among subjects.

This paper reviews the rationale for and development of a probabilistic atlas and reference system of the human brain derived from a large series of subjects, representative of the entire species, with retention of information about variability. The atlas includes structural as well as functional information. Such a project must take on the problems inherent in dealing with a variable biological structure and function but, when successful, provides a system that is realistic in its complexity, has defined accuracy and errors, and that, as a benefit, contributes new neurobiological information. Such a strategy will also spawn atlases of other species as well as human atlases of pathological conditions such as Alzheimer's

disease, autism, schizophrenia, multiple sclerosis (MS) and many others. These disease-specific atlases can then be used to demonstrate the natural history of disease progression and will find utility in clinical trials where experimental therapies can be examined for their impact on disease progression using an automated, objective and quantifiable reference atlas of the natural history of the disease state.

2. MOTIVATION FOR DEVELOPING A PROBABILISTIC HUMAN BRAIN ATLAS

The relationship between structure and function in the human brain, at either a macro- or microscopic level, is complex and poorly understood. Furthermore, we are not proposing to unravel this complexity with the data collected in the context of building this atlas. Rather, we will continue to develop a probabilistic framework in which appropriate datasets can be entered, across an ever-increasing number of modalities, between subjects, laboratories and experiments such that, in time, the aggregate data from populations will provide even greater (in both quality and quantity) insights into this important relationship. Our perspective on brain function is typically equated with the methods available to measure it. For the tomographic brain imaging techniques, the results produce macroscopic estimates of where gross functional changes (typically of a haemodynamic nature) are occurring. Electromagnetic techniques can provide direct information about when these events occur, and indirect information about where. The development of a probabilistic reference system and atlas for the human brain simply provides the framework in which to place these ever-accumulating data sets in a fashion that allows them to be related to one another and that begins to provide insights into the relationship between micro- and macroscopic structure and function.

(a) *Growth of neuroscience and lost opportunities.*

The growth of neuroscience in the last 25 years has been extraordinary. Annually, over 20 000 individuals attend the meeting of the Society for Neuroscience in the USA. At that meeting, over 40 000 papers have been presented in the last three years. Brain mapping and neuroimaging have witnessed a similar exponential rise in interest, output and productivity, although on a smaller scale. Throughout the neuroscience community, there is a general frustration with the volume of data that is generated and its relative inaccessibility in forms other than narrative text. Consider, for example, that over 13 000 Society for Neuroscience abstracts are published in hard copy and electronically each year. Faced with such a staggering volume of information, the individual neuroscientist typically retreats to his or her small scientific niche, resulting in ever-increasing specialization and isolation within the field.

At the same time, funding for neuroscience research has a limited return on its investment in that only a small fraction of raw data that is collected through such funds is analysed fully, and far less is interpreted and published. Even when published, narrative formats require arduous comparisons across experiments, methods and species.

If there were a system that provided a logical and organized means by which to maintain data from

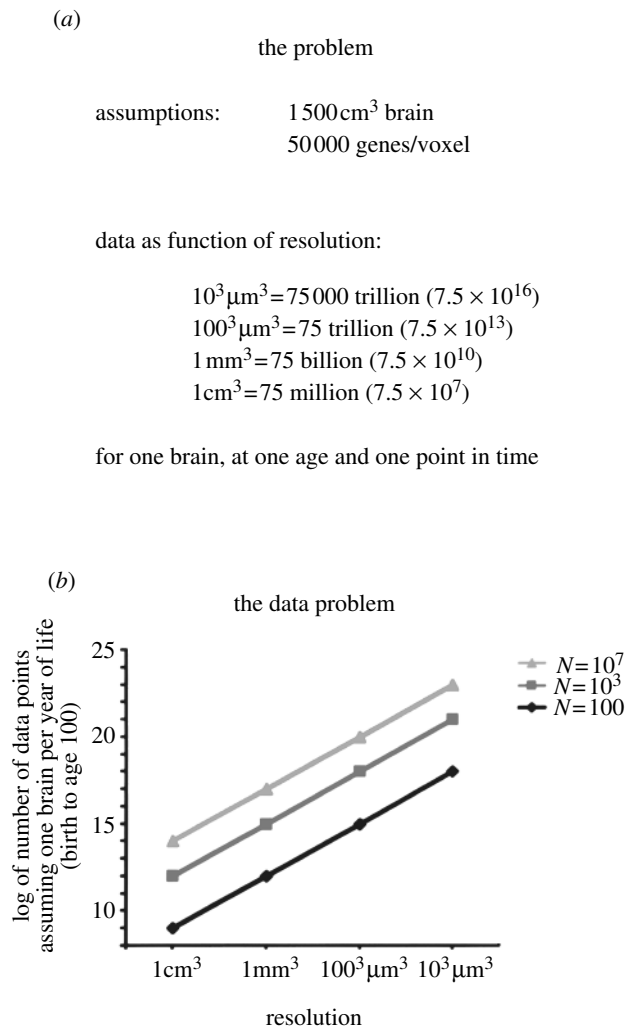


Figure 1. The magnitude of neuroinformatics data for the human brain. While there are a number of assumptions in this illustration, the orders of magnitude are realistic and enormous. They depict what would be involved in developing an organized data structure that combines location in the human brain with gene expression maps. (a) This example assumes that approximately 50 000 genes may be expressed in any three-dimensional region (voxel) of the brain at any given time during development. The typical human male brain is 1500 cm³ in volume. Based on the spatial resolution used to determine gene expression (ranging from 1 cm³ to 10³ μm³), the number of data points ranges from 75 million to 75 thousand trillion. Keep in mind that this is what is required to manage the data from just one brain at a given point in time. (b) If one uses these same assumptions and the range of resolutions noted in (a), the range of data magnitudes for a series of brains collected across a population with a representative of each age from birth to age 100 years old results in dataset magnitudes that range from 10⁹ to 10²³. These truly astronomical orders of magnitude will require innovative, practical neuroinformatic data structures that allow the referencing of such information as a function of both location and time.

meetings, individual experiments or the field as a whole, referenced to the anatomy of the brain, the species and the stage in development or duration of a pathological process, highly automated, content-based queries would vastly improve access, allow immediate comparisons among experiments and laboratories, provide a manage-

able format to assess new data at meetings or through periodical publications, provide for electronic experiments and hypothesis generation using the data of others to test theories, and greatly increasing the value of money spent on neuroscience research. Such an outcome requires sophisticated neuroinformatics tools, dedicated scientists committed to the successful completion of such a project in a practical fashion, and a paradigm shift in the sociology of neuroscientists with regard to information sharing (Koslow 2000). Nevertheless, the benefits of such an approach are enormous on their own, and of even greater value if one extrapolates the current situation to even greater numbers of neuroscientists and datasets in the future.

(b) Data suffocation

As the quality of neuroscientific data improves, so too does its magnitude. As spatial resolution in imaging data changes by one order of magnitude in one dimension, the volume of data points increases by a factor of 1000. *In vivo* imaging instruments are now routinely capable of producing 1 mm³ resolution elements whereas microscopic and ultrastructural studies achieve spatial resolutions 1000–100 000 times better. If one considers that 50 000–75 000 genes code for proteins of relevance to the human nervous system at some point during the life span, one can see the impact of assaying and storing such information across a range of spatial resolutions, as demonstrated in figure 1. Current genomic technology, and future advances in it, make feasible the ability to generate vast amounts of genetic information. All of these data are in search of an organizational home referenced to the location of the sample in neuroanatomical terms and the time frame of the sample as a function of the development of the organism. Once again, the brain's architecture becomes the most appropriate and intuitively sensible structure in which to organize such data so as to optimize correlations between biologically related datasets. While the remainder of this article focuses on imaging data of the human brain, it is important to note the magnitude of the information management problem for neuroscience as a whole.

(c) Data integration

To demonstrate the practical uses of the probabilistic reference system, an example is taken from actual experience, namely, an experiment performed by Watson and colleagues (Watson *et al.* 1993) to identify the visual motion area of the human brain (i.e. V5 or MT) (Ungerleider & Desimone 1986) using relative cerebral blood flow (CBF) (Mazziotta *et al.* 1985; Fox & Mintun 1989) measured with positron emission tomography (PET) (figure 2). In the experiment, each subject had multiple PET–CBF studies in two states, the first viewing a stationary pattern of targets and the second with the targets moving. The significant difference between the datasets collected in these two states (Friston *et al.* 1991) was then superimposed on magnetic resonance imaging (MRI) data using the automated image registration (AIR) algorithm (Woods *et al.* 1992, 1993). The result of this experiment demonstrated consistent bilateral activations of the dorsolateral, inferior occipital cortex in each subject. Furthermore, a consistent relationship between




Figure 2. Illustrative example of human visual area V5. (a) Four separate subjects demonstrating bilateral CBF-PET activations of V5 superimposed on their respective structural MRI studies (Watson *et al.* 1993). Note the consistent relationship seen between the activated site (red) and the ascending limb of the inferior temporal sulcus that also coincides with the cortical region (seen in (b)) identified by Flechsig (1920) as being myelinated at birth. (c) Patient studied by Zihl *et al.* (1983, 1991) with damage to the V5 area resulting in a selective disturbance of visual motion perception.

the site of increased blood flow and the frequently observed ascending limb of the inferior temporal sulcus (Ono *et al.* 1990) was found (figure 2a). Because the investigators were knowledgeable about temporo-occipital anatomy and physiology, they recognized that this location had also been identified by Flechsig (Flechsig 1920) as a portion of the human cerebral cortex (Flechsig Field 16) that is myelinated at birth (Bailey & von Bonin 1951) (figure 2b). These observations have been repeatedly confirmed by independent laboratories demonstrating the human V5 (Zeki *et al.* 1991) areas as a frequently detectable and robust functional landmark at the temporo-occipital junction (Dumoulin *et al.* 2000).

Now envision this experiment performed using neuroinformatics tools previously developed or proposed for the probabilistic reference system. Prior to performing the V5 PET experiment, each subject would perform a 'functional reference battery' of tasks, thereby providing functional landmarks throughout the brain. Following the experiment, anatomical warping and segmentation tools would be used to segment and label the anatomical regions of the brain automatically for each subject. Alignment and registration by functional landmarks would show the effects of functional registration on macroscopic anatomy. Functional alignment and registration using the V5 activation sites would automatically demonstrate the consistent relationship between that functional region and the ascending limb of the inferior temporal gyrus. Furthermore, it would

quantitate, in probabilistic terms, the spatial relationships between the sulcal-gyral anatomy and the functionally activated zone across subjects. Differences in responses could be related to demographic, clinical and genotypic information (Bartley *et al.* 1997; Zilles *et al.* 1997), if these were collected as part of the experiment, and related to population data already available in the 4-dimensional database. Cyto- and chemoarchitectural data, as it begins to populate the database, would be available for automated reference with regard to this cortical zone (Clark & Miklossy 1990; Rademacher *et al.* 1993). Time-series data from electroencephalography (EEG) or magnetoencephalography (MEG) would show the temporal relationships of this region to others (Dale *et al.* 1999; Ahlfors *et al.* 1999). Lesion data could also be accessed if such datasets had been added as an attribute (Zihl *et al.* 1991) (figure 2c). This is in contrast to the current situation where activated cortical regions are identified and one must laboriously search the literature to try to identify, qualitatively, in experiments with different characteristics, qualities and attributes, the regions of the brain that are of experimental interest for a given neuroscientific question.

3. STRATEGY AND RATIONALE

(a) Overall concept

The goal of the International Consortium for Brain Mapping (ICBM) is to develop a voxel-based, probabilistic atlas of the human brain from a large sample of normal individuals, aged eighteen to ninety years, with a wide ethnic and racial distribution. The dataset is designed to contain a substantial amount of demographic information describing the subjects' background, family history, habits, diet and many other features. In addition, clinical and behavioural evaluations include neurological examinations, psychiatric screening, handedness scores and neuropsychological tasks. One cubic millimetre multi-spectral MRI studies including T1-, T2- and proton density-weighted pulse sequences are obtained consistently. A subset of subjects also have functional imaging using a standardized battery of tasks and employing functional MRI, PET and event-related potentials. DNA samples will be acquired from 5800 of the subjects and made available for genotyping.

From an organizational point of view, eight laboratories in seven countries on three continents participate in the core data collection and analysis. These sites were selected because of their expertise in brain imaging, capacity to perform a large number of studies in a consistent fashion, and the fact that most sites had different imaging devices and computer platforms, thereby requiring the consortium to solve problems of interoperability and data differences from different acquisition devices.

It was decided early in the planning for the programme that in situations where the optimal solution to a given problem (e.g. data analysis pathway, visualization scheme, etc.) was not known, each laboratory would independently try to solve these problems. Once a laboratory-specific solution was obtained, appropriate algorithms would be distributed to consortium participants and evaluated. Ultimately, these algorithms were sent to outside laboratories for independent evaluation and comparison with methods developed by non-consortium groups. In each case, the

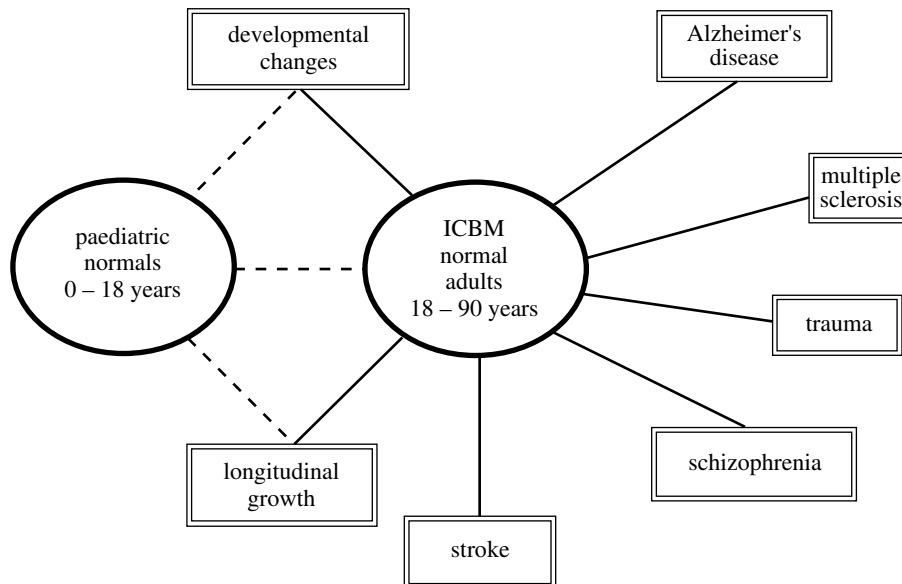


Figure 3. The normal adult ICBM atlas includes individuals between the ages of 18 and 90 years old, as indicated in the centre of this figure. As a result of the practices, principles and methods developed through this core project, a number of other atlases and databases have been spawned. Most logical was the development of a paediatric database for structural and functional brain anatomy. This is indicated on the left side of the figure. Through the use of datasets in both of these atlases as well as the methods developed in the consortium programme, it has been and will be possible to perform interesting cross-sectional and longitudinal studies across age ranges. A number of disease-specific atlases have been developed and are discussed in the text. A number of these are listed in this figure. It is anticipated that there will be a continuing growth in the number of different disease-specific atlases and that these atlases will all become related to the normal ICBM probabilistic atlas through a common focus on neuroanatomy and a four-dimensional data structure.

optimal strategy was then incorporated into the final approach used by the consortium. This was a 'real-world' situation designed to produce the optimal result through competition. As each successful component of these competitions emerged, it was incorporated into the overall ICBM strategy for data analysis, visualization and distribution. Thus, while each laboratory developed an independent strategy for processing data, the consortium as a whole made the commitment to a unified, centralized strategy for the pooled results, thereby resulting in a single atlas rather than a federation of atlases. The latter would result in inconsistencies in data analysis and confounding factors for users of the atlas in the long run.

The principles, practices and tools developed through the ICBM consortium have also spawned a series of other atlas projects on different populations (figure 3). Probabilistic atlases for children (i.e. from birth to eighteen years of age) and disease states (e.g. Alzheimer's disease, traumatic brain injury, MS, autism, schizophrenia, stuttering, cerebral infarction) are under development. These population- and disease-specific atlases have been developed for different reasons but employ similar principles and many of the same tools used for the normal adult brain atlas described here.

We also consider a part of this project to be the development of a reference system. The atlas will describe brain structure and function in three spatial domains and a temporal one referenced to the age of the subjects. Attributes (e.g. blood flow, receptor density, behaviours inducing blood flow changes at specific sites, signs and symptoms associated with lesions at specific sites, literature references) are then superimposed on the basic atlas.

As such, the atlas becomes the architectural framework for the reference system, the former being grounded in the four physical dimensions and the latter being extensible, based on the interests and datasets available to consortium participants and future users.

(b) *Probabilistic*

Since there is no single, unique representation for the human brain that is representative of the entire species, its variance must be captured in an appropriate framework. The framework that we have chosen is a probabilistic one in which the inter-subject variability is captured as a multi-dimensional distribution. These probabilities can change if subpopulations are sampled because of the shifting distributions. The probabilistic approach was relatively new to neuroanatomical thinking when we first proposed it in 1992. The only previous related strategies had to do with post-mortem analyses that reported distributions for structure sizes and dimensions for certain select regions of the brain (Filimonoff 1932). In recent years, the probabilistic strategy has been more widely used (Roland & Zilles 1994, 1996, 1998; Mazziotta *et al.* 1995*a,b*) and many probabilistic atlases are now being developed for such species as the monkey and the mouse.

For many psychiatric and behavioural problems, the relationship between structural abnormalities and disease is not straightforward. For example, while many studies have demonstrated that schizophrenics are more likely to have certain structural abnormalities shown by MRI (e.g. reduced brain size, ventricular enlargement, altered grey matter density in association cortex), none of these abnormalities is sufficiently distinctive or specific to make

MRI scanning a useful procedure in the routine clinical diagnosis of schizophrenia. Similar situations prevail for behavioural disorders such as dyslexia and autism.

In a recent paper, Leonard *et al.* (1999), employed a novel approach for using MRI to identify patients with schizophrenia. Rather than utilizing a single measure (e.g. ventricular volume), they used a linear discriminant function analysis of ten different anatomical measures derived using data from 37 schizophrenic patients and 33 normal controls. The controls had been recruited from among hospital staff, with balancing of age, handedness quotient and parental socio-economic status between the groups. However, the groups were not balanced for ethnic background or intelligence quotient (IQ)—and in fact, the control group had significantly higher verbal and performance IQs than the schizophrenic patients. Using their derived discriminant function, Leonard *et al.* (1999) were able to classify 79% of the controls and 76% of the schizophrenic patients correctly into the proper diagnostic category. Furthermore, they found that these 10 anatomical measures accounted for a significant part of the variance in measured full scale IQ in the schizophrenics but not in the normal controls.

This study illustrates a common concern that arises in comparisons of different groups of subjects with one another in a setting where many different anatomical measures are employed without clearly defined *a priori* hypotheses. Real differences between the groups among uncontrolled variables may be the true cause for differences that end up being attributed to the disease itself. For example, it is possible that the schizophrenic patients studied had a markedly different ethnic background from the hospital employees who were recruited as controls. It is also possible that the anatomical differences detected are, in fact, characteristic of any group of normal subjects having IQs in the range observed in the patients. Having a large, readily accessible, probabilistic database of normal subjects would make it much easier for investigators doing this type of research to evaluate the pertinence of uncontrolled variables. For Leonard's study, their linear discriminant function could be applied to a set of normal cases to determine whether the false positive rate for diagnosing schizophrenia varies as a function of ethnic background or IQ. Demonstration that it does not would bolster the notion that the anatomical findings really are a specific reflection of the schizophrenic disease process.

Ideally, the potential correlation between brain structure and behavioural measures would be explored even before recruiting subjects into this type of study. An investigator seeking to replicate Leonard's findings might use data selected from the large normal dataset to explore possible normal correlations between the ten anatomical measures and any of the many behavioural measures already recorded in the database. The large number of subjects in the database would allow even the extreme tails of the normal distribution to be explored for correlations that might not be evident in a small group of subjects clustered near the mean. Any behavioural measures that showed strong correlations could then be carefully matched between patient and control groups. To the extent that such matching is not feasible, quantitative results from the large normal dataset could be used to

estimate and discount the effects of group differences in behavioural measures on the dependent anatomical measures. Identifying and controlling for relevant variables that correlate with brain anatomy would be an essential element in establishing the credibility of derived, quantitative MRI measures that cannot be validated by simply having an experienced radiologist inspect a scan using his or her own internalized database of experience to define normality.

(c) *Neuroanatomy is the language of neuroscience*

(i) *Many nomenclatures*

The basic language of neuroscience is neuroanatomy. However, as in any global topic, many languages and dialects exist. Analogous to air traffic control systems, the ultimate solution to the development of a useable brain atlas requires location references expressed as coordinates and a common language to express them (for air traffic control, it is the English language). In developing the probabilistic atlas, it was our intention to be able to accommodate multiple languages and meanings. As such, it was important to build a hierarchical nomenclature system in which aliases could be referenced and the boundaries to which they referred adjusted, based on the language selected. This resulted in the requirement for a nomenclature editing system (BRAINTREE, see §3c(ii) below) and an approach that ultimately allows translation from one neuroanatomical language to another without the requirement to force all investigators to use a single, arbitrarily chosen language. It remains clear that the final solution does require a coordinate-based approach devoid of many of the ambiguities associated with qualitative naming of structures.

(ii) *BRAINTREE*

We have developed a system that provides a graphical relationship between anatomical nomenclature and its relationship with the structure or system to which it belongs. To link this nomenclature to a three-dimensional space from the atlas, BRAINTREE relies on a two-coordinate bounding box for each of the nodes, producing a defined region of three-dimensional space that entirely encompasses the named structure. The user can select a structure on the basis of its standard nomenclature and have its coordinates passed on to standard display or measurement tools. Hence, the BRAINTREE program provides a facile interface between an editable hierarchical nomenclature system and the indexable three-dimensional coordinate space. Furthermore, the nomenclature can easily be extended to include the myriad of aliases that are common in neuroanatomy or even relate the structural names that provide an association between species (Toga *et al.* 1996).

(d) *Use of a large population*

It is clear that the use of large populations is an essential requirement in the development of an atlas that is intended to capture the variance in structure and function of the human brain. Such a large population can be newly acquired, as was done in this project, or could be the result of pooling smaller studies to produce a meta-database. The latter approach was rejected because, after examining reports in the literature of smaller sample size

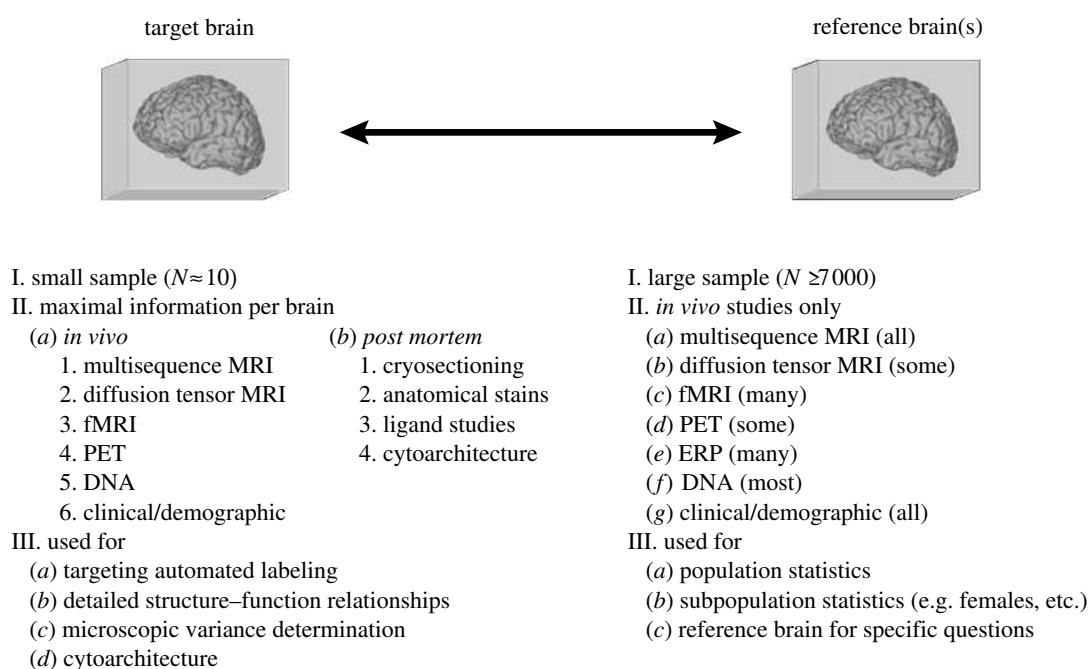


Figure 4. Target versus reference brains. Fundamental to the development of the probabilistic atlas is the concept that there are two types of datasets that are required for a comprehensive system. The first is a target brain or brains. The criteria for this resource are listed in the figure. Ideally, single individuals who are normal and studied during life will have the complete complement of *in vivo* imaging studies performed. After death, these same individuals would be studied again using appropriate imaging studies, such as MRI, and then detailed cytoarchitectural and chemoarchitectural analyses of their brains would result in a very information-rich ultimate dataset. Because it is unlikely that such a situation will occur (although we have studied a few individuals ante- and post-mortem), typically, the post-mortem and *in vivo* studies are obtained on separate individuals. These target brains are used to understand microscopic and macroscopic structure–function relationships better but also to take a newly acquired dataset and warp it to match the target, thereby picking up the anatomical labels for each structure on a voxel-by-voxel basis. When the new study is back-transformed to its original shape and configuration, all structures in the brain will have the appropriate labels with an error rate defined by the confidence limits of the warping algorithm and other factors discussed in the text. Reference brains are also described in the figure. These represent large populations of *in vivo* studies where population statistics about variance for structure and function can be obtained. Reference brains can also be sampled to provide subpopulations appropriate to a unique set of descriptors. The size of the sample set for the reference brains will dictate whether subpopulations of sufficient size will be available for specific projects.

projects, it was clear that there was such a wide range of methodological and strategic differences among these studies as to make their pooling difficult, if not impossible. Technical issues such as voxel size, slice thickness, scanning parameters and many others would cause difficulties in any attempt to produce a homogeneous final product. The same can be said of subject selection and description. As might be expected, a wide range of criteria were used in selecting subject populations, including the definition of normality. Screening tests, demographic and background information as well as neurological and psychiatric examinations vary from study to study, adding to the incompatibility of the pooled results. If one also includes functional information, the situation is far worse. Since brain function is obtained by having subjects perform tasks, any slight variation in the task presentation, psychophysics or the strategy employed by the subject in performing the task will cause unpredictable differences among experiments, thereby adding methodological variance in the pooled data and confounding the final product. Thus, it was decided to collect a sample of a large number of subjects prospectively for which these confounding factors could be controlled.

Given the need to have much larger populations of subjects than had previously been available, the current

programme is now intended to include 7000 normal subjects obtained from geographical locations as disparate as Japan and Scandinavia and spanning the age range from 18 to 90 years. Special efforts have been made to obtain a wide range of racial and ethnic diversity. In addition, 342 twin pairs (half mono- and half dizygotic) are also part of this sample. The dataset for each subject includes a detailed historical description of medical, developmental, psychological, educational and other demographic features. In addition, behavioural data, including neurological, neuropsychological and neuropsychiatric examinations, are part of the dataset. In 5800 subjects DNA samples are being collected, stored and made available for genotyping. This large sample size allows the opportunity to provide realistic estimates about the variance of structure and function for brain regions, the relationships between structure and function at macro- and microscopic levels, and true phenotype–genotype–behavioural comparisons. The large sample size also increases statistical power in making such inferences about the population or when the atlas is used as a comparison sample for investigations involving other groups, be they normal or pathological. Lastly, as the sample size increases, the opportunity to select subpopulations of meaningful size also increases.

(e) Target and reference brains

A fundamental concept of the consortium's project was to distinguish between target and reference brains (figure 4). We have defined the target brain to be the dataset, derived from one or, ideally, a few individuals, that has the richest collection of data available. Theoretically, this would be the brain of a normal individual studied with *in vivo*, high resolution, structural and functional imaging and then, after death, having detailed post-mortem analysis including cyto- and chemoarchitecture. If a series of such brains could be studied, then a probabilistic target brain would emerge. Given the high resolution of the post-mortem data, target brains would be the most informative with regard to anatomical and chemical localizations. While we have studied a few individuals (all elderly) who had both *in vivo* macroscopic brain imaging and—through the UCLA Willard Body Program—post-mortem cryosectioning, we typically do not have both *in vivo* and post-mortem datasets of the same individual. As such, a synthesis of this information into an optimized target brain has been the practical solution, to date.

In contradistinction to the target brain, reference brains are derived from large populations of subjects typically through *in vivo* imaging of structure and function. These datasets provide information about variance in the population for both structure and function, but at a three-dimensional spatial resolution that is three orders of magnitude lower than the target brains.

Target and reference brains are used for different purposes. Target brains are, as the name implies, the target to which an unlabelled dataset can be warped. The unlabelled dataset then picks up the anatomical, functional or other attributes of each voxel. Once it is back-transformed to its original shape, the new dataset will have the appropriate anatomical and functional labels for all brain regions. A certain percentage of these labels will be erroneous based on imperfections of the warping system, an incomplete understanding of the anatomy of homologous brain regions between subjects and errors in the primary labelling of the target brain. Reference brains provide data about distributions of brain regions and can be divided into subpopulations for specific purposes. Reference brains give estimates of anatomical and functional regions in a population of individuals and, as such, can be used to determine confidence limits when a new dataset falls outside the range of normality or expected variance for a given population. Taken together, these two tools provide important but very different vehicles for analysing existing or new datasets with regard to brain structure and function.

(f) Function

It is important to emphasize from the outset that our motivation for studying functional landmarks in this project is analogous to the motivation for studying structural anatomical landmarks. Specifically, the ICBM atlas will use functional landmarks to augment atlas methods that are currently based primarily on macroscopic structural anatomy in the same way that these anatomical methods now augment atlas methods that were previously based on stereotaxis with simple proportional scaling. An important distinction must be made between

functional imaging to answer neuroscience questions, which is not proposed here, and functional imaging to serve as a neuroinformatics tool, which is our intent. Whereas neuroscience functional imaging studies currently use individual macroscopic anatomical landmarks to define a common neuroinformatics framework for comparing and combining data from different subjects, we anticipate that future neuroscience functional imaging studies will complement this macroscopic anatomy with functional anatomical landmarks, identified in each individual through a selected battery of neuroinformatics tasks. In general, we expect that the neuroinformatics tasks and the functional landmarks that they produce may be completely unrelated to the tasks constituting the primary focus of the neuroscientific investigation. It is our objective to develop and validate tasks that are well suited to producing functional landmarks. These tasks will be the first of what we expect to be an ever growing library of tasks that will constitute a functional reference battery (FRB). The FRB will be used to develop a new generation of brain atlases through novel warping techniques that move beyond macroscopic anatomy and into the realm of functional and cytoarchitectonic similarities as the fundamental basis for homologous mapping of one brain to another.

The major theoretical and practical issue in identifying homologous brain structures, and in warping strategies designed to compare brains within a population of subjects, is a critical issue with regard to both three-dimensional and surface geometries and representations. In this project we have based all aspects of the atlas development on three-dimensional, voxel-based strategies. Nevertheless, this neither obviates nor limits one's capacity to address special issues related to cortical surface topology. In fact, a significant fraction of the programme has been focused on the development of appropriate cortical surface extraction and cortical interface (e.g. grey–white interface) identifications. It is important to understand the appropriate constraints that must be imposed to preserve cortical surface topology for both the cerebrum and the cerebellum (Felleman & Van Essen 1991; Van Essen & Drury 1997; Van Essen *et al.* 1998; Fischl *et al.* 1999).

To understand the motivation to identify functional landmarks, it is important to understand the differences and similarities between functional landmarks and anatomical landmarks with respect to meeting the objectives of neuroinformatics research. Neuroanatomy and, specifically, neuroanatomical landmarks have been the basis that formed the framework for indexing neuroscience information from a number of specific sources collected across spatial scales. Explicit in the plan was the notion that the atlas system would need to continue to adapt in an iterative fashion to accommodate improvements in spatial scale and in the models used to map data into a single neuroanatomical framework. A self-critical evaluation of the methodologies used for structural atlas alone reveal areas in need of extension:

- (i) *Macroscopic landmarks from structural MRI studies provide a suboptimal basis for appropriate mapping of individual anatomy into a unified neuroinformatics framework*

Three independent lines of research serve to demonstrate the difficulties of relying exclusively on macroscopic

anatomical landmarks as a neuroinformatics framework. The first evidence comes from the significant progress made in warping three-dimensional anatomical data to match a target template. In the absence of brain pathology, it is computationally feasible to use high order, nonlinear warps to generate a one-to-one correspondence between brains, even while requiring perfect alignment of unambiguous cortical anatomical features such as the crests of gyri and the depths of sulci. However, even with the inclusion of such anatomical constraints, these mappings are not unique. Mechanical properties such as viscosity or elasticity must be ascribed to the brain tissues to find a solution that is optimal from the standpoint of those presumed mechanical properties (Christensen *et al.* 1993; Davatzikos 1997; Schormann *et al.* 1996). In the absence of more restrictive constraints or independent external standards, different solutions can all lead to equally good (from the standpoint of visual inspection or image similarity criteria) but mutually inconsistent answers, indicating that with macroscopic anatomical data alone, the mapping problem is substantially under-constrained. While it might be a valid computer science goal to identify the transformation that perfectly maps one brain onto another, while minimizing some intuitively appealing quantity, this is not necessarily the best neuroinformatics goal. A more appropriate goal from a neuroinformatics standpoint is to maximize the genuine homology of points that are brought into correspondence by the transformation. Functional landmarks will provide additional constraints on inter-subject warping that will help to meet this important neuroinformatics goal.

Various criteria can be used to define homology, and conflicts between macroscopic homologies and microscopic cytoarchitectonic homologies (Rademacher *et al.* 1993) constitute the second line of research demonstrating the problems of relying solely on macroscopic structural anatomy. Recent post-mortem cyto- and chemoarchitectonic studies have shown that even some sulcal and gyral features that were once thought to be almost perfectly correlated with nearby cytoarchitectonic boundaries are in fact only approximately correlated (Zilles *et al.* 1997; Geyer *et al.* 1997, 1999, 2000; Amunts *et al.* 1999, 2000). It is our explicit bias that homologies based on function and cytoarchitectonics are more fundamental to neuroscience, and hence to its informatics, than homologies based on sulcal and gyral anatomy.

The third line of research that highlights the difficulties of an informatics framework that is based solely on structural anatomy comes from the rapidly expanding field of functional magnetic resonance imaging (fMRI). A decade ago, functional imaging with PET was of sufficiently low resolution that atlases based on the simple proportionality of the original Talairach system were adequate to assure that homologous activation sites would overlap from subject to subject and that the resulting group results would be interpreted as consistent across laboratories. Subsequent improvements in PET image resolution have justified the adoption of the more sophisticated techniques based on structural MRI scanning and MRI–PET coregistration that are in widespread use today (Woods *et al.* 1993). The high resolution possible with fMRI, and the fact that statistically significant responses are readily identified in fMRI data from a

single subject, demand much more accurate mapping of homologous landmarks from every individual subject and threaten to make methods that rely solely on macroscopic anatomy obsolete. Ideally, a neuroinformatics framework should seek to stay a step ahead of such developments. Functional links may also be of particular value for patient populations where normal function may persist even in the presence of substantial anatomical distortions. Providing the necessary link between global and local anatomy is a problem that will require new tools, new approaches and new population data acquired specifically for that purpose.

(ii) *Cytoarchitectonic studies provide an insufficient basis for quantifying relevant inter-subject variability in the population*

As mentioned above, cytoarchitectonic studies in a small number of subjects can be extremely powerful in demonstrating the potential range of inter-subject variability—hence, our motivation to begin to incorporate such data into the ICBM atlas. However, the collection of cytoarchitectonic data is extremely demanding in terms of time and resources. These realities make it unlikely that reliable population estimates of the variability between structural and functional anatomy will be quantified for many brain regions any time soon using these techniques. Since cytoarchitectonics cannot be identified *in vivo*, such data may help to define general rules (e.g. a given cytoarchitectonic field is most likely to be located at position X in women and at position Y in men), but will not help to identify the individualized exceptions to such rules. In contrast, functional imaging is well suited to population based studies and functional imaging can be applied routinely to living individuals. To a first approximation, functional landmarks can be viewed as an *in vivo* proxy for cytoarchitectonic landmarks. It should be explicitly stated that it is not our primary intent to equate a given functional landmark with a given cytoarchitectonic region. Indeed, it is clear that a one-to-one relationship will sometimes not exist, since functional subdivisions are present as maps within some cytoarchitectonic areas (e.g. M1 and V1) and since adjacent, functionally correlated, areas can be distinguished cytoarchitectonically. Rather, we view cytoarchitectonic anatomy and functional anatomy as intrinsically intertwined features that reveal an underlying pattern of brain organization that provides an optimal framework for neuroscience research and neuroinformatics challenges. By warping brains in a way that brings homologous functional landmarks into concordance, we expect simultaneously to bring nearby cytoarchitectonic regions into better superimposition, even if we do not explicitly know the identities of the cytoarchitectonic regions or even the locations of their boundaries (figure 5). Capturing the unique spatial information represented by functional landmarks is an important front for neuroinformatics research—one that will provide routine, direct access to this fundamentally important level of brain organization.

Properties of good and informative functional landmarks

The minimal attributes of a good functional landmark are that it be unambiguously detectable in individuals and that the variability in its location within individuals be small. ‘Small’ is a relative term, and contexts for making this judgement will be explicitly defined below,

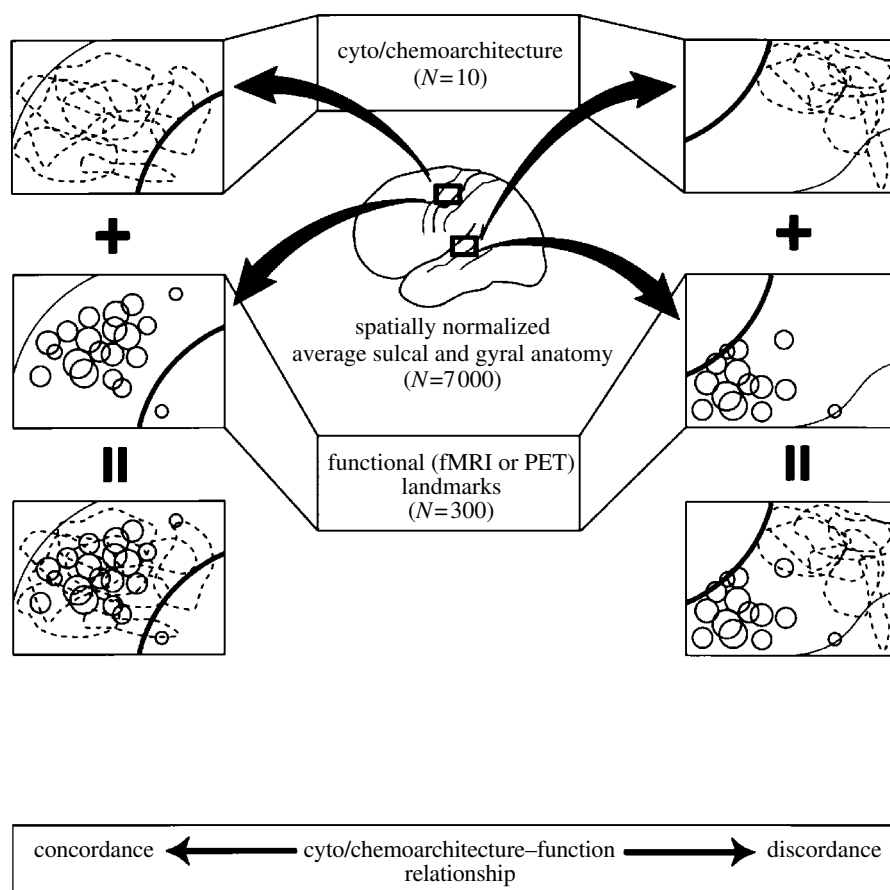


Figure 5. Hypothetical example of concordant and discordant relationships between macroscopic function, derived from fMRI or PET, and microscopic cyto- or chemoarchitecture from post-mortem specimens. First, all datasets are spatially normalized in probabilistic space at a macroscopic scale (e.g. gyri/sulci) then the relative positions of functional landmarks (centre row) are examined relative to the sites of cyto- or chemoarchitectural zones (top row). There is superimposition of the sites in the left panels (concordance) but not in the right panels (discordance). With systems and sets of data such as those proposed in the ICBM atlas, it will be possible to add an ever-increasing body of this type of information, leading to structure–function relationship insights throughout the cortex, deep nuclei, cerebellum and brainstem.

along with specific consideration of how ‘landmarks’ can be extracted from functional images. Additional desirable attributes of functional landmarks and the tasks that produce them are listed in table 1. We make an important conceptual distinction here between a ‘good’ functional landmark and an ‘informative’ functional landmark. In order also to be considered informative, a good functional landmark should provide unique information that could not have been determined purely on macroscopic anatomical grounds. For example, a functional task for identifying primary visual cortex might not prove to be particularly informative since human cytoarchitectonic data indicates that striate cortex consistently maps to the calcarine fissure (Polyak 1957) (though some variability is present, see review by Aine *et al.* 1996). In contrast, a good functional landmark in a frontal region, where gyral anatomy is quite variable, might be highly informative. Caution is generally indicated when trying to predict in advance which functional landmarks will ultimately prove informative since detailed studies of the relationship between cytoarchitectonic and macroscopic anatomy are still relatively rare, as are data comparing functional and structural anatomy. Indeed, with careful study, some traditional assignments of functional areas to specific sulcal or gyral locations are proving to be less reliable

than previously expected (Aine *et al.* 1996; Zilles *et al.* 1997). Likewise, data from some functional studies have identified previously unsuspected function–structure correlations. A good example of this latter situation comes from studies of putative human area V5, where substantial variability in Talairach coordinate location across subjects turned out to be largely explained by a highly consistent relationship between V5 (defined functionally) and the intersection of the ascending limb of the inferior temporal sulcus and the lateral occipital sulcus (figure 2a) (Watson *et al.* 1993). Because of the difficulty in predicting which landmarks will be informative, we have primarily focused our attention on identifying tasks that produce good functional landmarks and identifying these landmarks in a representative population. These data are being evaluated to determine how informative the landmarks actually are and to look for currently unrecognized structure–function correlations.

Those who are primarily involved in functional imaging neuroscience (as opposed to neuroinformatics) research may be surprised that our criteria for a good functional landmark do not include that the landmark should have a consistent location across subjects. When trying to answer neuroscience questions, there are situations where variability across subjects is undesirable—one

Table 1. Criteria for evaluating functional landmarks

essential criteria	good functional landmarks		informative functional landmarks
	pragmatic criteria	desirable criteria	criteria
universally (or nearly universally) identifiable in individual subjects without ambiguity	location insensitive to environmental variation (e.g. background noise levels, room lighting)	location independent of imaging modality	must meet essential criteria for good functional landmarks
location in individuals stable with repeated testing	location insensitive to educational background, native language, gender, etc. subject performance verifiable or irrelevant minimal opportunity for diverse cognitive strategies tasks simple enough to be applicable to cognitively impaired patient populations or children	underlying physiology understood identifiable simultaneously with many other landmarks produced by a single task and control	must provide unique spatial information not predictable from macroscopic anatomy

hopes that a functional task will produce responses at a highly consistent, anatomically standardized location across subjects so that overlapping regions of response will increase statistical significance and so that the consistency of location will increase confidence that the areas seen in each individual are truly homologous. It is, therefore, perhaps counterintuitive that exactly the opposite situation applies to functional landmarks to be used for neuroinformatics. A functional landmark that is always present in the exact same location in every subject, when using current methods of anatomical standardization, is assured to be uninformative, providing only redundant information that could have been derived from the anatomical data alone.

The neuroinformatics goal here is to use functional landmarks to provide a new source of valid, independent anatomical information that cannot be detected using macroscopic anatomy and to use this information to improve the homologous mapping of different subjects to one another or to an atlas. The result should be better mapping from one subject to another that will serve to improve local homology, a goal that should prove advantageous when subsequently analysing neuroscience functional imaging data in these same subjects. Two major and one minor assumption are implicit in this line of reasoning and need to be explicitly stated: (i) despite the variation in location, it is critical that the functional landmarks that are identified in each subject are truly homologous; (ii) methodological variability in establishing the location of the functional landmark within each subject must be small when compared with the true anatomical variability in the standardized location of the landmark across subjects; and notable, though less important; (iii) some preservation of local topology is assumed, so that establishing the location of a functional landmark will indeed improve the homologous mapping of nearby brain regions.

An important implication of the last two assumptions is that the value of a functional landmark will vary: (i) depending on the amount of within-subject variability

(more variability decreases its value); (ii) depending on the amount of local inter-subject variability (more variability increases its value); and (iii) depending on its proximity to the nearby regions where better mapping is desired (greater proximity increases its value). If the goal is to improve mapping throughout the brain, numerous functional landmarks may be needed, whereas local mapping may be improved with just one strategically placed functional landmark. The value of proximity raises an important consideration: in functional neuroscientific imaging experiments, why bother to use the locations of established functional landmarks that may be unrelated to the task of interest rather than simply using the locations produced by the primary task itself? There are at least two good answers to this question: first, unless landmarks produced by the primary task have been determined to be good landmarks (implying considerable prior investigation), the resulting mapping may actually lead to less reliable homologous mappings than anatomical data alone; and second, statistical models for evaluating group significance would be invalidated by such a procedure unless separate trials were used for mapping and for addressing the primary neuroscience question. Consequently, appropriate use of the landmarks produced by the primary task being investigated would require that these landmarks be validated and used in exactly the same way as any other nearby functional landmark. The use of landmarks will also depend, in part, on the brain region(s) of interest for a given experiment and the interests of the investigator.

(g) *Analysis strategy*

At the outset of this project, it was unclear what the optimal analysis strategy would be for both the structural and functional aspects of the programme. Given the large number of subjects, each with multi-spectral MRI datasets and many with functional imaging studies as well, it was clear that the tools to be developed would have to function in an automated, or at least semi-automated,

fashion to be feasible. Furthermore, reliable automaticity would be a general benefit to the brain imaging field, given the labour-intensive aspects of manual image editing. It was also clear that certain steps would be required to process data in what we have called an ICBM 'analysis pipeline'. These steps include:

- (i) screening data for obviously incomplete or artefact-laden studies and rejecting them;
- (ii) intensity normalization in three dimensions for each pulse sequence;
- (iii) alignment and registration across pulse sequences and studies within a given subject;
- (iv) tissue classification (i.e. grey and white matter, cerebrospinal fluid, other);
- (v) 'scalping' whereby extracranial structures are removed;
- (vi) spatial normalization of each subject to a target where anatomical labels can be obtained automatically;
- (vii) surface feature extraction;
- (viii) visualization.

Given this sequence of tasks, it was unclear, in most cases, what the optimal solution for each would be. Rather than making an *a priori* decision and have all consortium members work to achieve it, an alternative approach was chosen. It was decided that each of the primary laboratories in the consortium would work to solve each step in the analysis pipeline independently and in parallel. These laboratory-specific algorithms would then be locally optimized. Once a given laboratory was satisfied with the performance and documentation of their approach, it would be distributed to the other participating laboratories for alpha testing. If an algorithm failed to perform adequately, or was awkward to use because of hardware platform incompatibilities or other factors, it was rejected. Those algorithms that performed well across consortium laboratories were ultimately sent to an independent group (David Rottenberg, Stephen Strother and colleagues at the University of Minnesota) for beta testing. This independent testing included not only the ICBM algorithms for a given module in the analysis pipeline but also any other algorithms that could be identified worldwide that purported to perform the same functions. During beta testing, algorithms were evaluated with simulated as well as real datasets selected by the beta test laboratory and evaluated for documentation, ease of installation, computation time, accuracy and precision. The results of these evaluations were then published (Strother *et al.* 1994; Arnold *et al.* 2001). The winners of this competition were then selected for the ICBM analysis pipeline (figure 6) and will be the basis for the mass data analysis of all datasets. While it was decided that it was important to analyse all 7000 studies in a consistent manner, so that users would know the methodology, algorithms and versions of the algorithms from which the results were derived, this in no way precluded individual laboratories in the ICBM consortium or elsewhere from using their own strategies for data analysis on the original datasets which are provided through digital libraries (see §3i(i)). This strategy has been successful in that it established an internal competition whereby the best solution emerged rather than using

an *a priori* and hypothetical prediction that might have fallen far short of the optimal outcome.

(h) *Visualization*

Similar to the approach chosen for analysis, it was decided to keep an open mind as to how to present the data developed by the consortium. Given the probabilistic nature of the resultant data, the decision is not straightforward and has not yet been fully resolved. It may well be that the optimal solution is to select many avenues and that users of this system choose for themselves. Approaches that flatten (Carman *et al.* 1995; Felleman & Van Essen 1991; Van Essen *et al.* 1997, 1998) or inflate (Dale *et al.* 1999; Fischl *et al.* 1999) the cortex have been proposed and well described. As a visualization tool, these strategies allow cortical anatomy to be seen in its entirety at the expense of the more familiar, three-dimensional appearance of the brain. It is important to note here that visualization, simply as a tool to view the data, must be distinguished from the use of these tools to identify homologies between regions in different brains or different species. In this context we consider these strategies only as a visualization tool as our approach to homology identification was described earlier with regard to macro- and microscopic structure–function considerations. Each visualization strategy has its benefits and limitations. The traditional three-dimensional view of the brain in its natural state obviates the ability to see brain regions hidden in folded cortex or deep structures without providing tools for translucency or sectioning. Flattening or inflating the surfaces will produce areas of compression and expansion that alter the data from its original state but make all surface regions visible. Providing all of these avenues will allow the user to choose among them given a specific purpose. The user can choose whether the benefits and insights provided by a given visualization strategy outweigh the disadvantages or artefacts induced by the visualization scheme.

(i) *Database*

(i) *Digital libraries*

In addition to the derived data organized in the databases described above, digital libraries and data warehouses (figure 7) of complete datasets will also be provided through the ICBM project to the neuroimaging community. These datasets include those with 'raw' data (i.e. complete, three-dimensional, multi-spectral MRI structural studies of individual subjects), 'scalped' (i.e. extracranial structures removed) datasets, and intensity normalized, 'scalped' datasets. Access to such information may allow investigators to obtain normal control data for neuroimaging experiments or to test various methods for image analysis and display without the requirement to acquire original data on their own. Most problematic will be the distribution of the 'raw' dataset, since the potential for compromising subject confidentiality is an issue. Since the experimental subject's face could be reconstructed from the raw datasets, one strategy would be to alter or eliminate facial structures from the dataset prior to distribution.

(ii) *Four-dimensional*

There currently exists no comprehensive database for the storage of complete, individual subject, neuroimaging

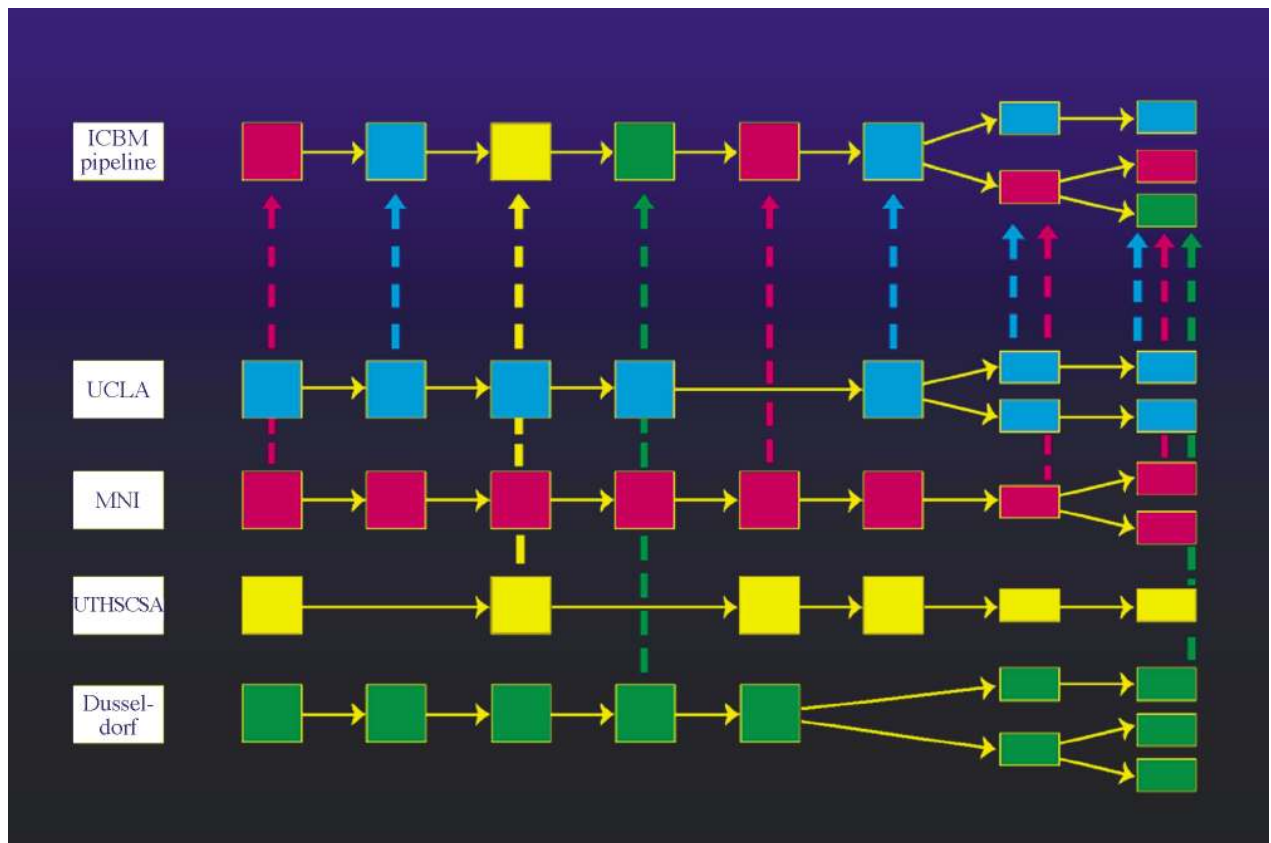


Figure 6. ICBM analysis pipeline. Since it was impossible to predict what specific algorithm or mathematical strategy would be optimal for each step in the analysis of structural and functional data collected for subjects in this consortium project, the original core laboratories elected to each develop independent strategies for each step. Once complete and tested within a given laboratory, they were distributed among consortium participants for alpha testing. After the consortium members were satisfied with the performance at this phase, all consortium-developed algorithms were delivered to an independent laboratory (that was, not part of the consortium), for beta testing. The beta testing included not only the ICBM-developed algorithms but also any other algorithms identifiable worldwide that purported to perform the same function. The best (see § 3g) algorithm was then selected for incorporation in the ICBM pipeline. All data in the final atlas will be processed through this unified, single pathway. The bottom four boxes (white) in the left column represent consortium sites contributing algorithms to the pipeline.

datasets for the human brain that is both electronically accessible and efficient in its interactions with neuroscientists. This reduces the value of both clinical and research funds spent on the acquisition of these important and interesting studies. The physical world is organized in four dimensions and, thus, forms a logical and comprehensive organizational framework for the ICBM database. Plans anticipate the future inclusion of time-series data from dynamic, functional data acquisition methods such as fMRI, EEG and MEG, requiring the fourth dimension. It is expected that spatio-temporal and purely temporal patterns of brain activity will constitute functional entities and markers of their own. These can be used for the following purposes.

- (i) Since function will be defined in the future by brain locations and timing of activity, the probabilistic reference will incorporate temporal and spatio-temporal brain activity information.
- (ii) Spatio-temporal and temporal functional markers will be used for most of the same purposes described for the spatial functional markers (fMRI, PET) in this paper: warping; correlations across subjects; an additional source of information in calculating population distributions these can be

used, for example, to inform studies with small populations, etc.

- (iii) Temporal and spatio-temporal information will be used to correlate brain activity across subjects in the temporal dimension.
- (iv) They can also be used as priors for brain source estimation methods, and their probability distributions can be used for Bayesian procedures in EEG and MEG brain source localization (Schmidt *et al.* 1999).

With this data structure, queries-by-content tools and strategies are being developed. These tools will allow users of the database to submit a query in the form of actual data (e.g. a two-dimensional image of a portion of the brain or a three-dimensional block of data) and ask the database to search for matches using wavelet-based techniques that have previously been demonstrated to be successful for two-dimensional internet searches of graphic material (Wang *et al.* 1997).

The expansion of these approaches to three and, eventually, four dimensions will be an important neuroinformatics milestone that will find uses far beyond the applications in this consortium.

Furthermore, a system organized in this fashion, and the tools associated with it, will allow for efficient,




Figure 7. Digital libraries and data warehouses. We envision two outputs of the probabilistic atlas to the neuroimaging and neuroscientific communities. The first is termed 'digital libraries'. In these libraries we will catalogue complete MRI studies from large numbers of subjects. They will be held in three forms: raw data (with the face corrupted for confidentiality reasons); intensity normalized images; and 'scalped' intensity normalized images where extracranial tissues have been removed. A user would be able to select a certain list of descriptors (e.g. females, aged 25–30, right-handed) and will be told how many subjects and datasets fulfil this requirement. These datasets can then be transferred (either electronically or through magnetic media, depending on the size) to the requesting investigator. Data warehouses provide database interactions where probabilistic structures and high-level queries can be obtained by interacting with the entire population of 7000 studies or user-selected subpopulations. Such queries can be representational or symbolic and may ultimately be object-oriented. These queries may involve the imaging, clinical, behavioural or genetic data.

convenient and comprehensive access by neuroscience clients to the ever growing data in the ICBM probabilistic reference system. The goal is not to develop physiological models of brain function, neural connectivity and other important neurobiological questions. But these exciting opportunities will be more easily achieved by providing a system of database interactions and structure for modelers, neuroimagers and neuroscientists in general. We envision that, once established and populated with data, the probabilistic reference system that is organized in this fashion will allow for 'electronic' hypothesis generation and experimentation using previously collected, well-described and effectively organized data.

(iii) *Attributes*

It is conceptually important to understand that the database architecture, while organized in four dimensions, can have a very high number of attributes all referenced to these basic four dimensions. These additional attributes need not be specified at the time of establishing the data sample or the dataset. Some can be derived and others can be added at a later point through further examination of the original subjects (e.g. longitudinal studies, other methodologies) or by further analysis of existing data (e.g. genotyping of stored DNA samples).

The most difficult challenge to the actual organization of such a database is the scaling and referencing of data across major spatial or temporal domains. While originally developed to have a fundamental spatial unit of

1 mm³ resolution, there is no reason why microscopic and ultrastructural information cannot appropriately populate the individual 1 mm³ voxels of the macroscopic dataset. The same can be said of temporal information, but the exact manner of binning of time-series information will require judicious attention to the types of queries anticipated of such datasets.

(iv) *Central versus distributed*

A business metaphor is appropriate here. Fledgling industries rarely do well when trying to establish standards, means of communication and interoperability methods that are designed to result in a reliable and durable outcome for a given community. Examples abound, including: telecommunications, aviation, electronics, meteorology and others. In most of these cases, a well-designed centralized approach established both the problems and the solutions that later led to deregulated, decentralized systems that were linked by regulatory groups, industrial standards and meta-databases. Similarly, in the burgeoning field of neuroinformatics, an initial centralized approach appears both desirable and manageable. It allows for a straightforward and easily monitored means of distributing datasets on a continuous basis. A centralized approach can also monitor the required submission of attributes derived from the datasets back into the database as a measure of successful, reciprocal sharing of data and results (Bloom 1996; Pennisi 1999). Finally, in order to even attempt such a

project, there must exist a critical mass of data analysis tools, organization and reputation to make participation attractive psychologically and sociologically. The common goal and ultimate result must also be sufficiently valuable to the contributing sites to make participation compelling. The results of participation must be worth more than the sum of the individual parts.

(j) *Real world*

The ICBM consortium has always maintained a 'real world' environment such that the participating sites use different equipment, software and protocols reflecting a microcosm of the larger neuroscience, neuroimaging and neuroinformatics communities and forcing the development of solutions to problems through flexible, compatible systems rather than rigid standards, protocols and equipment requirements. The significance of this feature is that the products are not platform-, institution- or protocol-specific.

(i) *Interoperability*

Interoperability was an important concern early in the development of the ICBM atlas. So important was the requirement to develop interoperable tools and datasets that a conscious decision was made to deliberately utilize imaging instruments, computing hardware and file formats that differed among the participating sites. This forced certain principles and rules to be utilized in the development of software and the exchange of data, the goal being accessibility of any ultimate end user to all of these products. The psychology and sociology of any advanced research field is to develop home-made tools and to maintain intralaboratory file structures. The experience in the ICBM consortium was no different. As such, we developed translators that would allow datasets to be transferred among sites with an agreed upon file format (MINC, Neelin *et al.* 1998) but that was translated into the 'home' file format upon receipt at any of the participating sites. A similar strategy was used for algorithms. This simplistic approach has worked quite well, allowing a relatively seamless exchange of information.

(ii) *Quality control*

If the ICBM atlas is to be a growing resource, tools that have been developed, thus far, will ultimately be open to the entire neuroscientific community for the future additions of datasets. How then will we assure the quality of data from investigators? Having pondered and debated this question for many years and having examined the approaches used by other fields, the simple answer is that we cannot assure a certain level of quality control in a completely open data exchange program. Not only is this impractical but it may also lead to the erroneous exclusion of data that might someday be deemed valuable. If there were some filter on the input of data, what would the review process be? How can we predict how tomorrow's observations will be judged by today's standards? We cannot. Furthermore, in a practical sense, such an approach would immediately become backlogged with datasets awaiting 'review' by some 'panel of experts' whose opinions might change as time and experience progresses. What we can provide, however, is a system by which users of such datasets can select their own level of

confidence about the populations or results that they sample. For example, a user might request all information about a certain region of the brain for a given demographic population of subjects. Most of these data would be of high quality and reliably collected but some of them would undoubtedly include experimental, methodological and other errors. Nevertheless, it would give the user a complete picture of all of the information available about their query. At the other end of the spectrum, consider a user who is interested in only the most accurate information about a given site in the brain for a certain population. That user could request data that was only obtained from the results of peer-reviewed, published and independently reproduced data collections. Thus, just as the datasets can be filtered using demographic, anatomical or clinical criteria, they can also be filtered and queried by confidence level. 'Let the user beware' is the only rational approach to developing such a system.

4. METHODS AND RESULTS

(a) *MRI*

(i) *Basic principles*

Multi-spectral anatomical MRI data for the ICBM project were acquired using optimized protocols matched as closely as possible across the different scanner manufacturers and field strengths (3.0 T GE (Milwaukee, WI), 1.5 T Philips (Bothell, WA), and 2 T Elscint (Israel)). The protocol design goals were to achieve whole head 1mm isotropic T1-weighted image volumes and whole head $1 \times 1 \times 2$ mm T2- and PD-weighted volumes.

(ii) *Averaging*

The accuracy of brain atlases is constrained by the resolution and signal gathering powers of available imaging equipment. In an attempt to circumvent these limitations, and to produce a high resolution *in vivo* human neuroanatomy, we investigated the usefulness of intra-subject registration for *post hoc* MR signal averaging (Holmes *et al.* 1998). Twenty-seven high resolution (7×0.78 mm³ and 20×1.0 mm³) T1-weighted MRI volumes were acquired from a single subject, along with twelve double-echo T2/PD weighted volumes. These volumes were automatically registered to a common stereotaxic space in which they were subsampled and intensity averaged. The resulting images were examined for anatomical quality and usefulness for other analytical techniques.

The quality of the resulting images from the combination of as few as five T1 volumes was visibly enhanced. The signal-to-noise ratio was expected to increase as the root of the number of contributing scans, to 5.2 for an n of 27. The improvement in the $n=27$ average was great enough that fine anatomical details, such as thalamic subnuclei and the grey bridges between the caudate and putamen, became sharply defined. The grey-white matter boundaries were also enhanced, as was the visibility of any finer structure that was surrounded by tissue of varying T1 intensity. The T2 and PD average images were also of higher quality than single scans but the improvement was not as dramatic as that of the T1 volumes. Overall, the enhanced signal in the averaged images resulted in higher quality anatomical images with improved results for other post-processing techniques. The high quality of the enhanced images permits novel uses of the data and extends the possibilities for *in vivo* human neuroanatomical explorations. *Post hoc* registration and




Figure 8. Coronal image demonstrating tritiated muscarinic receptors from one hemisphere of a cryosectioned brain and demonstrating the anatomical detail that such chemoarchitectural maps can provide. When serial sections are obtained and stained for a wide range of receptors, anatomical features and gene expression maps, a tremendous wealth of information is available for comparison with sites of functional activation obtained using *in vivo* techniques and macroscopic brain structure (gyri, sulci, deep nuclei, white matter tracts). Having a probabilistic strategy for relating these different types of anatomies will provide previously unavailable insights about the relationship of structure and function on both microscopic and macroscopic levels for the human brain and, by analogy, for the brains of other species (see figures 5 and 9). The analysis of the regional and laminar distribution patterns of transmitter receptors is a powerful tool for revealing the architectonic organization of the human cerebral cortex. We succeeded in preparing extra-large serial cryostat sections through an unfixed and deep-frozen human hemisphere. Neighbouring sections were incubated with tritiated ligands for the demonstration of 15 different receptors of all classical transmitter systems. The distribution of [³H]oxotremorine-M binding to cholinergic muscarinic M2 receptors is shown here as an example. Even a cursory inspection of a colour-coded receptor autoradiograph permits the distinction of numerous borders of cortical areas and subcortical nuclei by localized changes in receptor density and regional/laminar patterns. For example, the M2 receptor subtype clearly labels the primary sensory cortices (at the level of the section shown in the figure, e.g. the primary somatosensory area BA3b and the primary auditory area BA41) by very high receptor densities sharply restricted to both areas. The different receptors allow the multi-modal molecular characterization of each area or nucleus by the so-called receptor fingerprint typing. A receptor fingerprint of a brain region consists of a polar plot based on the mean density of each receptor in the same architectonical unit (area, nucleus, layer, module, striosome, etc.). The following areas and nuclei could be delineated in the present example: (i) cingulate cortex; (ii) motor cortex; (iii) primary somatosensory cortex (BA3b); (iv) inferior parietal cortex;

averaging of MRI scans is a robust method for the enhancement of MR images. There is a significant reduction in noise in averaged images that reveals previously unobservable structure. The high quality of the resulting images opens the door to other forms of post-processing and suggested even further applications. The method itself is very straightforward and can easily be employed.

(iii) *Quality control issues*

A major challenge in the oversight and ultimate outcome of a large multi-site project involving investigators from around the world is to maintain quality assurance and control. While this entails innumerable administrative issues, not appropriate for this review, one area is of interest. Since the decision was made to collaborate among sites having different scanner manufacturers, the problem of attempting to determine an appropriate means of calibration and quality control among these sites became an issue. Typically, manufacturers will provide a phantom specific for their instrument that identifies aberrations in radio frequency, field homogeneities and geometrical distortions. Such phantoms are typically not optimal for use among different instruments in the MRI field. In addition, we experienced numerous problems trying to transfer phantoms and related calibration materials across international borders in an efficient manner. This led to the concept of the living or 'smart' phantom. The smart phantom is an individual (or group of individuals) who physically travel from site to site at regular intervals, bringing with them a physical phantom as well. Thus, the same brain(s) is scanned at each site serially over time. This provides a number of advantages that we had not been able to achieve using physical phantoms alone. First, it provides a convenient way of moving the physical phantom from site to site and ensuring that it is scanned on time and in a correct manner. Second, it provides a stable and realistic dataset that can be used for calibration among different participating sites. Third, there are certain liaison activities associated with this process, in that the 'smart phantom' interacts directly with those individuals most specifically involved in scanning test subjects, an opportunity that affords interactions, suggestions, trouble-shooting and fact-finding. Fourth, this strategy provides a convenient way of calibrating a new or upgraded instrument, an inevitable event in a study of this duration, prior to the scanning of test subjects.

(b) *Post-mortem cryosectioned material*

(i) *Data acquisition*

Mapping the human brain and its functions requires a comprehensive anatomical framework. This reasoning dictated the need in our consortium to obtain high resolution, digital, whole brain, post-mortem datasets. The fact that recent advances in anatomical digital imaging techniques now permit unrestricted visualization in multiple cut planes, and three-dimensional regional or subregional analyses when appropriate primary datasets are available (Spitzer & Whitlock 1992; Wertheim 1989), made this approach feasible. Digital representations also offer the opportunity for morphometric comparisons and sophisticated mapping between anatomical and metabolic

(v) insular cortex; (vi) primary auditory cortex (BA41); (vii) non-primary auditory cortex; (viii) inferior temporal association cortex; (ix) entorhinal cortex; (x) mediodorsal thalamic nucleus; and (xi) putamen (K. Zilles, A. Toga, N. Palomero-Gallagher and J. Mazziotta, unpublished data).

imaging modalities (Payne & Toga 1990; Toga & Arnica-Sulze 1987). The primary source data for human brain atlas must include not only very fine spatial detail but also image colour and texture to convey the subtle characteristics that make it possible to distinguish subnuclear and laminar differences. Furthermore, the incorporation of an appropriate spatial coordinate system is critical as a framework for inter-subject morphometrics. High-resolution anatomical datasets serve as references for the accurate interpretation of clinical data from the PET, computed tomography (CT) and MRI modalities as well as the mapping of transmitters, their receptors (figure 8) and other regional biological characteristics.

Thus, we have designed a system of histological and digital processing protocols for the acquisition of high resolution, digital imagery from post-mortem cryosectioned whole human brain and head for computer-based, three-dimensional representation and visualization (Cannestra *et al.* 1997; Toga *et al.* 1997). High-resolution (1024^2 pixel) serial images can be captured directly from a cryoplaned blockface using an integrated colour digital camera and fibre-optic illumination system mounted over a modified cryomacrotome. The system can process tissue treated in a variety of ways, including fixed, fresh, frozen or otherwise prepared for sectioning at micrometre increments. Sometimes it is desirable to section the tissue while still *in situ*. Specimens frozen and sectioned with the cranium intact preserve brain spatial relationships and anatomical bony landmarks. Colour preservation is superior in unfixed tissue but unfixed heads were incompatible with decalcification and cryoprotection procedures. Thus, section collection from such specimens was complicated by bone fragmentation. Collection of 1024^2 images from whole brains results in a spatial resolution of $200\ \mu\text{m}/\text{pixel}$ in a 1–3 gigabyte data space. Even higher three-dimensional spatial resolution is possible by primary image capture of selected regions such as hippocampus or brainstem or by using higher resolution cameras. Discrete registration errors can be corrected using image processing strategies such as cross-correlative and other algorithmic approaches. Datasets are amenable to resampling in multiple planes as well as scaling and transpositioning into standard coordinate systems. These methods enable quantitative measurements for comparison between subjects or to atlas data. These techniques allow visualization and measurement at resolutions far higher than those available through other *in vivo* imaging technologies, and provide greatly enhanced contrast for delineation of neuroanatomical structures, pathways and subregions.

The use of cryosectioned anatomical images as a gold standard for mapping the human brain requires a complete understanding of the assumptions and errors introduced by this method. While there are several obvious advantages to using these data as a reference for other tomographic and *in vivo* mappings, their collection requires sophisticated instrumentation and representative post-mortem material. Spatial resolution, the inclusion of bony anatomy, full colour, blockface reference for histologically stained sections and the resulting registered three-dimensional volumetric datasets are important aspects of this method. Nevertheless, cryosectioning approaches, like all others, introduce distortion during acquisition and processing. Sources of errors include post-mortem brain changes and artefacts associated with tissue handling. A major source of error is related to specimen preparation prior to sectioning. Removal of the cranium and subsequent brain deformation, perfusion protocols or freezing altered the spatial configuration of the dataset.

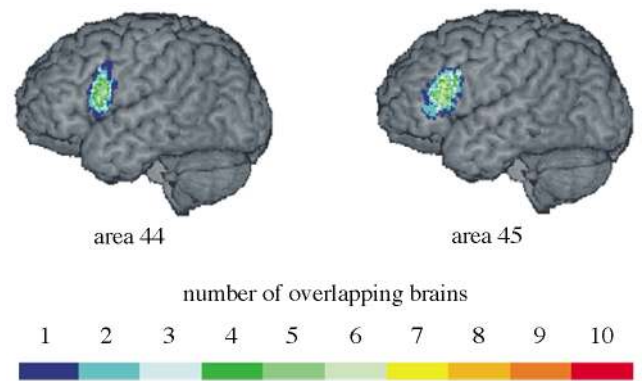


Figure 9. Location and extent of Broca's region (Brodmann areas 44 and 45). Areas 44 and 45 as defined in serial coronal sections of an individual brain after three-dimensional reconstruction; lateral views of the left hemisphere. Probability maps of Broca's region, based on microscopic analysis of ten human brains can be referenced, also in a probabilistic fashion, to functional activation sites associated with the functions of Broca's area using the multimodality probabilistic atlas strategy. The overlap of individual post-mortem brains is colour-coded for each voxel of the reference brain (colour bar), e.g. seven out of ten brains overlapped in the yellow-marked voxels. Left in the image is left in the brain.

While three-dimensional data at this resolution is difficult to acquire, it is necessary for careful studies of morphometric variability and the generation of digital comprehensive neuro-anatomical atlases (Thompson *et al.* 1995). Ultimately, what is needed is the combined use of cryosectioned data as the source of higher resolution raw and stained anatomy spatially referenced to an *in vivo* electronically, acquired dataset such as those obtained with MRI.

(ii) *Cyto- and chemoarchitecture*

A major effort in this project is to obtain cyto- and chemoarchitectural data from post-mortem brains to enter into the probabilistic database for comparison with *in vivo* studies. An example of this approach is described for Broca's area. The putative anatomical correlates of Broca's speech region, i.e. Brodmann's areas 44 and 45 (Brodmann 1909), are of considerable interest in functional imaging studies of language. It is a long-standing matter of discussion whether or not anatomical features are associated with the functional lateralization of speech (Galaburda 1980; Hayes & Lewis 1995, 1996; Jacobs *et al.* 1993; Simonds & Scheibel 1989; Scheibel *et al.* 1985). Furthermore, the precise position and extent of both areas in stereotaxic space and their inter-subject variability still remain to be analysed, since Brodmann's delineation is highly schematic, not documented in sufficient detail and does not contain any statement about inter-subject variability.

We studied the cytoarchitecture of Brodmann's areas 44 and 45 in ten human post-mortem brains using cell body-stained (Merker 1983) $20\ \mu\text{m}$ thick serial sections through complete brains (Amunts *et al.* 1999). Cytoarchitectonic borders of both areas were defined using an observer-independent approach, which is based on the automated high resolution analysis of the packing density of cell bodies (grey level index, GLI) from the border between layers I and II of the cortex–white matter junction (Schleicher *et al.* 1999). These profiles are perpendicular to the cortical surface and define the laminar pattern of cell bodies. Thus, the profiles are a quantitative expression of




Figure 10. (a) Autosegmentation of structures. This image illustrates the first stage of autosegmentation once the brain has been spatially normalized. Lobes, gyri and some subcortical nuclei are labelled. This iterative process continues with increasing refinement. (b) Three-dimensional model with autosegmented ventricular system. This model shows an autosegmented ventricular system converted to a surface model enabling morphometric statistics to be calculated. This segmentation was a combination of tissue classification approaches (note the intensity gradient of CSF to grey (white matter is high) and template matching following spatial normalization).

the most important cytoarchitectonic feature. Multivariate statistical analysis was used for locating significant differences between the shapes of adjacent GLI profiles along the cortical extent. Those locations represent cytoarchitectonic borders. GLI profiles were also used for investigating inter-hemispheric differences in cytoarchitecture. Significant inter-hemispheric differences in cytoarchitecture (i.e. differences in GLI profiles between right and left areas) were found in both areas 44 and 45. Profiles obtained as internal controls from the neighbouring ventral premotor cortex did not show any lateralization.

The position of the borders of areas 44 and 45 with respect to sulci and gyri showed a high degree of inter-subject variability (figure 9). This concerned the sulcal pattern, i.e. the presence, course and depths of sulci, as well as the spatial relation of areal borders with these sulci. The position of a cytoarchitectonic border could vary by up to 1.5 cm with respect to the bottom of one and the same sulcus in different brains. Thus, sulci and gyri are not reliable and precise markers of cytoarchitectonic borders. Although there was a considerable inter-subject variability in volume of areas 44 and 45 ($n = 10$), area 44 was larger on the left than on the right side in all cases of our sample. We could not find any significant left–right differences in the volume of area 45.

The extent and position of areas 44 and 45 were analysed in the three-dimensional space of the standard reference brain of the European Computerized Human Brain Database (Roland & Zilles 1996) after the above described microstructural definition of the areal borders. MR imaging (3-D FLASH-scan, Siemens 1.5 T Magnet) was performed on post-mortem brains prior to histology. Corrections of deformations inevitably caused by the histological technique were performed by matching MRI and corresponding histological volumes (Schormann & Zilles 1997; Schormann *et al.* 1995). Brain volumes were finally transformed to the spatial format of the reference brain. For both steps, a movement model for large deformations was applied (Schormann *et al.* 1997; Schormann *et al.* 1996; Schormann & Zilles 1998). The superimposition of individual cytoarchitectonic areas in the standard reference format resulted in probability maps (figure 9). These maps quantitatively describe the degree of inter-subject variability in extent and position of both areas. They serve as a basis for topographical interpretations of functional imaging data obtained in PET and fMRI experiments (Amunts *et al.* 1998).

The observed inter-subject variability in the extent and cytoarchitecture of Broca's region has to be considered when correlating data of functional imaging studies with the underlying cortical structures. Inter-hemispheric differences in the volume of area 44 and in the cytoarchitecture of both areas may contribute to functional lateralization which is associated with Broca's region.

(c) *Warping and segmentation strategies*

(i) *Segmentation*

Manual voxel segmentation and labelling

We developed a general image analysis package, DISPLAY, which provides a wide range of capabilities for: (i) interactive three-dimensional exploration of image volumes using simultaneous orthogonal planes and surface-rendered representations; (ii) manual labelling of image voxels; (iii) archival/recall of labelled three-dimensional objects such as brain regions, pathological masses, tissue class maps, etc.; and (iv) morphological operations such as the dilate/erode/open/close primitives. DISPLAY has become a standard utility within the ICBM consortium, and elsewhere, for labelling brain regions (Evans *et al.* 1996; Paus *et al.* 1996a,b; Penhune *et al.* 1996). However, the use of manual tools for labelling large numbers of MRI datasets is prohibitively time-consuming and subject to inter-rater variability. We have therefore developed a series of algorithms for automated image segmentation (figure 10).

Correction for three-dimensional intensity non-uniformity—N3

A major problem for automated MRI image segmentation is the slowly varying change in signal intensity over the image, caused principally by non-uniformities in the radio-

frequency field. Apparent signal from any one tissue type is therefore different from one brain area to another, confusing automated segmentation algorithms that assume constant signal for one tissue type. We have developed a fully automated three-dimensional technique for inhomogeneity correction. The method maximizes the entropy of the intensity histogram to maximize its structure. The effect of inhomogeneity is modelled as a convolution histogram by a blurring kernel and the effective kernel can be estimated and deconvolved by iterative entropy maximization. The method is applicable to any pulse sequence, field strength and scanner (Sled *et al.* 1997, 1998). In the previously described competition among algorithms, the N3 approach of Sled *et al.* (1997, 1998) proved superior (Arnold *et al.* 2001) and has been selected for the ICBM data analysis pipeline.

Tissue classification—intensity normalized stereotaxic environment for classification of tissues (INSECT)

We have developed a series of algorithms for tissue classification (Kamber *et al.* 1992, 1995; Zijdenbos *et al.* 1996). They are used for automatically processing multi-spectral (T1-, T2-, proton density (PD)-weighted) datasets from large numbers of subjects, known as INSECT (intensity normalized stereotaxic environment for classification of tissues). All data are corrected for field inhomogeneity (Sled *et al.* 1998), inter-slice normalization and inter-subject intensity normalization. Stereotaxic transformation is then performed (Collins *et al.* 1994) and an artificial neural network classifier identifies grey/white/cerebrospinal fluid (CSF) tissue types (Zijdenbos *et al.* 1996; Evans *et al.* 1997).

These same tissue classification strategies are equally applicable for population analysis of patients with brain disorders and for tracking structural change over time, such as the progressive tissue atrophy that occurs in some degenerative diseases. We present one illustrative example of this approach, drawn from a clinical trial of a new treatment for MS which used MRI observations as a surrogate marker for disease activity (Zijdenbos *et al.* 1996). Multi-spectral MRI data were collected at 14 sites in North America from 460 patients with relapsing–remitting MS. A total of 1850 datasets were available, each consisting of T1-, T2- and PD-weighted volumes. After correction for MRI intensity inhomogeneity, interslice and intervolumetric intensity normalization, and stereotaxic transformation, the multi-spectral data were tissue classified to identify MS lesion voxels for each patient timepoint. Figure 11 shows a three-dimensional rendering of the probabilities for lesion distribution obtained from all datasets. This shows the most likely locations for MS lesions within a population and is a convenient way to distil a large amount of population data into a single entity. Tests of drug effect are reduced to testing for a significant group difference in the overall volume of this distribution above a given threshold when partitioned into drug and placebo groups. Tests for regional drug effects, i.e. changes in regional lesion probability, between groups become equivalent to the familiar test for cerebral blood flow (CBF) change between two stimulus conditions, using the same statistical models as are used for detection of functional changes in activation experiments with PET and fMRI (Evans *et al.* 1997). Importantly, this approach allows for rapid re-analysis of clinical trial data in response to: (i) new hypotheses; (ii) modification of the input data for discriminating lesion and non-lesion (types of image, noise filtering, image blurring); or (iii) modification of criteria for identifying lesions (threshold values for lesion probability, spatial constraints on lesion location, minimum voxels per lesion, etc.). This approach makes it feasible to extract the




Figure 11. Probabilistic atlas of multiple sclerosis (MS). This image is produced from 460 patients with MS derived from 5800 individual pulse sequences. The green areas demonstrate the probabilistic location of the actual MS plaques in the population referenced to the human ventricular system (red). This strategy gives a composite view of overall disease burden across a large population. Consider its use in a clinical trial where these patients were randomly assigned to treatment groups that included placebo, conventional therapy and experimental therapy. By following the group for some period of time and serially imaging them, an automated, quantifiable and objective measure of the relative effect of experimental therapy versus conventional therapy and the natural history of the disease could be obtained using MRI lesions as a surrogate marker of disease burden. Such strategies should lead to more efficient and cost-effective clinical trials.

considerable and usually untapped information available in large clinical trial image databases.

Regional parcellation—ANIMAL

Manual labelling of brain voxels is both time-consuming and subjective. We have developed an automated algorithm to perform this labelling in three dimensions (Collins *et al.* 1995). The ANIMAL algorithm (automated non-linear image matching and anatomical labelling) deforms one MRI volume to match another, previously labelled, MRI volume. It builds the three-dimensional, non-linear deformation field in a piecewise linear fashion, fitting cubical neighbourhoods in sequence. The algorithm is applied iteratively in a multi-scale hierarchy. At each step, image volumes are convolved with a three-dimensional Gaussian blurring kernel of successively smaller width (32, 16, 8, 4 and 2 mm full width at half maximum (FWHM)). Anatomical labels are defined in the new volume by interpolation from the original labels, via the spatial mapping of the three-dimensional deformation field.

(ii) *Warping strategies*

Atlases can be greatly improved if they are elastically deformable and can fit new image sets from incoming subjects. Local warping transformations (including local dilations, contractions and shearing) can adapt the shape of a digital atlas to reflect the anatomy of an individual subject, producing an individualized brain atlas. Introduced by Bajcsy and colleagues at the University of Pennsylvania (Broit 1981; Bajcsy & Kovacic 1989; Gee *et al.* 1993, 1995), this approach was adopted by the Karolinska

Brain Atlas Program (Seitz *et al.* 1990; Thurfjell *et al.* 1993; Ingvar *et al.* 1994), where warping transformations were applied to a digital cryosectioned atlas to adapt it to individual CT or MR data and co-registered functional scans.

Image warping algorithms, specifically designed to handle three-dimensional neuroanatomical data (Christensen *et al.* 1993; 1996; Collins *et al.* 1994, 1995; Thirion 1995; Rabbitt *et al.* 1995; Davatzikos, 1996; Thompson & Toga 1996; Bro-Nielsen & Gramkow 1996; Schormann *et al.* 1996, 1997; Schormann & Zilles 1998; Ashburner *et al.* 1997; Woods *et al.* 1998) can transfer all the information in a three-dimensional digital brain atlas onto the scan of any given subject, while respecting the intricate patterns of structural variation in their anatomy. These transformations must allow any segment of the atlas anatomy to grow, shrink, twist and rotate, to produce a transformation that encodes local differences in topography from one individual to another. Deformable atlases (Seitz *et al.* 1990; Evans *et al.* 1991; Miller *et al.* 1993; Gee *et al.* 1993; Christensen *et al.* 1993; Sandor & Leahy 1994; 1995; Rizzo *et al.* 1995) resulting from these transformations can carry three-dimensional maps of functional and vascular territories into the coordinate system of different subjects. The transformations can also be used to equate information on different tissue types, boundaries of cytoarchitectonic fields and their neurochemical composition (Amunts *et al.* 1998, 1999, 2000; Geyer *et al.* 1996, 1997, 1999, 2001).

Warping algorithms calculate a three-dimensional deformation field that can be used to non-linearly register one brain with another (or with a neuroanatomical atlas). The resultant deformation fields can be used subsequently to transfer physiological data from different individuals to a single anatomical template (Geyer *et al.* 1996; Larsson *et al.* 1999; Naito *et al.* 1999, 2000; Bodegård *et al.* 2000*a,b*). This enables functional data from different subjects to be compared and integrated in a context where confounding effects of anatomical shape differences are factored out. Non-linear registration algorithms, therefore, support the integration of multi-subject brain data in a stereotaxic framework, and are increasingly used in functional image analysis packages (Seitz *et al.* 1990; Friston *et al.* 1995).

Any successful warping transform for cross-subject registration of brain data must be high-dimensional, in order to accommodate fine anatomical variations (Christensen *et al.* 1996; Thompson & Toga 1998). This warping is required to bring the atlas anatomy into structural correspondence with the target scan at a very local level. Another difficulty arises from the fact that the topology and connectivity of the deforming atlas have to be maintained under these complex transforms. This is difficult to achieve in traditional image warping manipulations (Christensen *et al.* 1995). Physical continuum models of the deformation address these difficulties by considering the deforming atlas image to be embedded in a three-dimensional deformable medium, which can be either an elastic material or a viscous fluid (Schormann *et al.* 1996). The medium is subjected to certain distributed internal forces that reconfigure the medium and eventually lead the image to match the target. These forces can be based mathematically on the local intensity patterns in the datasets, with local forces designed to match image regions of similar intensity.

(iii) *Automated methods*

The inter-subject differences in the anatomy of the brain can be large, even after alignment, making anatomical segmentation inaccurate (Galaburda *et al.* 1978; Geschwind & Levitsky 1968; Gur *et al.* 1980; Steinmetz *et al.* 1991; Zilles *et al.* 1995,

1997). Without any perceived pathology, structures in the brain can differ in shape and size, as well as in relative orientation (Roland & Zilles 1994; Mazziotta *et al.* 1995*a,b*). Affine transformation is often insufficient for the labelling and segmentation of structures. Automated image registration algorithms can be used to align MR data with previously labelled and segmented brains by maximizing a measure of intensity similarity, such as three-dimensional cross-correlation (Collins *et al.* 1994), ratio image uniformity (Woods *et al.* 1992), or mutual information (Viola & Wells 1997; Wells *et al.* 1997). These techniques can be used in a non-linear fashion to obtain better results, but they still develop errors with small structures and in the borders of larger structures. The following steps have been used:

- (i) Each pulse sequence for each subject is intensity normalized within and between slices.
- (ii) Pulse sequences are aligned and registered within subjects (between pulse sequences) using AIR (Woods *et al.* 1992).
- (iii) The intensity normalized and aligned datasets from each subject are spatially normalized to our labelled target.
- (iv) Skull and scalp stripping is accomplished using the Leahy algorithm (Sandor & Leahy 1997).
- (v) Manual editing and segmentation is performed with SEG (UCLA) or DISPLAY (MNI).

To ensure accuracy, a previously labelled atlas is registered to the target brain via a non-linear technique that captures the desired structures in the region of interest (ROI) defined by a probability density of where the structure of interest lies. An iterative procedure is used by which the ROI in the atlas is registered to the ROI projected into the target brain via a high-dimensional warping technique that allows all segments of the anatomy to grow, shrink, twist and rotate. The ROI can then be refined to include greater detail and a closer approximation to the desired structure. The refined ROI in the atlas is again registered with the high-dimensional techniques to the target brain. This is repeated until the ROI is equal to the desired anatomical structure, within some allowed error estimate. In this way, a successive approximation, from lobar, to gyral, to subgyral, to nuclear resolution labels, can be achieved.

The registration techniques for the high-dimensional warps do not have to be limited to the previously mentioned intensity based techniques. Edges and surfaces can be automatically computed and used to determine boundaries of structures (Sandor & Leahy 1997; Lohmann 1998; Duta *et al.* 1999; LeGoualher *et al.* 1999; Zhou *et al.* 1998; Zhou & Toga 1999). These boundaries can be aligned and used to align the tissue that surrounds them via continuum mechanical techniques guiding the tissue flow (Thompson *et al.* 1996*a,b*).

Ultimately, it is the combination of different registration techniques in the proper order that archive a registration accurate enough to transfer the boundaries of one segmentation in the atlas to that of the target volume. When combined with the automated selection of a given atlas from a database of populations and the use of probabilistic information from the template associated with the class of interest, it will be possible to accurately label and segment any digital brain volume.

(d) *Surface methods*

(i) *Surface extraction*

Vast numbers of anatomical models can be stored in a population-based atlas (Thompson & Toga 1997, 2000*c*). These models provide detailed information on the three-dimensional

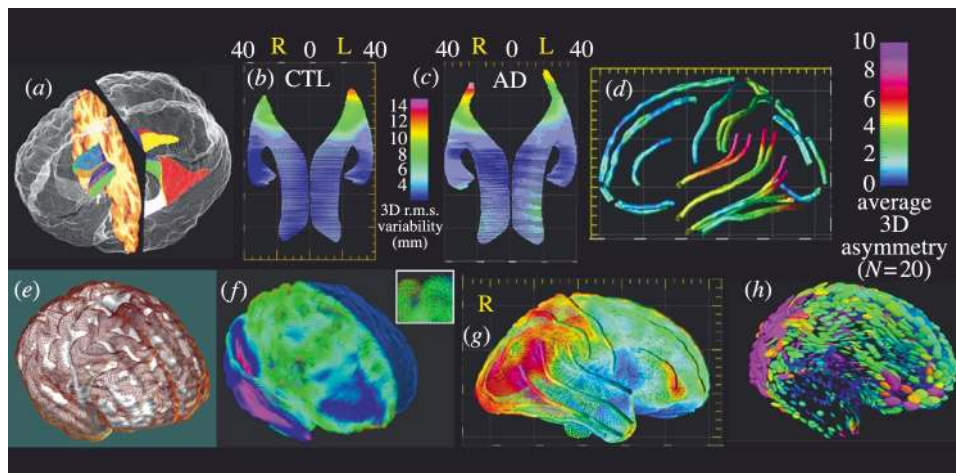


Figure 12. Surface models. Three-dimensional models can be created to represent major structural and functional interfaces in the brain. Panel (a) shows a model of the lateral ventricles, in which each element is a three-dimensional parametric surface mesh. Panels (b) and (c) show average ventricular models from a group of patients with Alzheimer's disease ($n=10$) and matched elderly controls ($n=10$). Note the larger ventricles in the patients and a prominent ventricular asymmetry (left larger than right). These features only emerge after averaging models for groups of subjects. Average population maps of cortical anatomy (d) reveal a clear asymmetry of the perisylvian cortex. Panel (e) shows an individual's cortex (brown mesh) overlaid on an average cortical model for a group. Differences in cortical patterns are encoded by computing a three-dimensional elastic deformation (f) (pink colours, large deformation) that reconfigures the average cortex into the shape of the individual, matching elements of the gyral pattern exactly. These deformation fields provide detailed information on individual deviations and can be averaged across subjects to create three-dimensional variability maps, demonstrating fundamental patterns of anatomical variability in the brain (g) (Thompson *et al.* 2000c). Tensor maps using colour ellipsoids (h) reveal the directions in which anatomical variation is greatest. The ellipsoids are more elongated in the directions in which structures tend to vary the most. Pink colours denote the largest variation while blue colours show the least. These statistical data can be used to detect patterns of abnormal anatomy in new subjects. Severe abnormality is detected (red colours) while corresponding regions in a matched elderly control subject are signalled as normal.

geometry of the brain and how it varies in a population. By averaging models across multiple subjects, subtle features of brain structure emerge that are obscured in an individual due to wide cross-subject differences in anatomy (Thompson *et al.* 2000b,c). These modelling approaches have recently uncovered striking patterns of disease-specific structural differences in Alzheimer's disease (Thompson *et al.* 1997, 1998, 2000b), schizophrenia (Narr *et al.* 2000) and fetal alcohol syndrome (Sowell *et al.* 2001), as well as strong linkages between patterns of cortical organization and age (Thompson *et al.* 2000a), gender (Thompson *et al.* 2000b), cognitive scores (Mega *et al.* 1997) and genotype (Le Goualher *et al.* 2001). To illustrate the approach, figure 12 shows a model of the lateral ventricles in which each element is represented by a three-dimensional surface mesh. These surface models can often be extracted automatically from image data, using recently developed algorithms based on deformable parametric surfaces (Thompson & Toga 1996; Thompson *et al.* 1996a,b; MacDonald 1998) or voxel-coding (Zhou & Toga 1999). Once an identical computational grid (or surface mesh) is imposed on the same structure in different subjects, an average anatomical model can be created for a group. This is done by averaging the three-dimensional coordinate locations of boundary points that correspond across subjects.

Figures 12b and 12c show average ventricular models from a group of patients with Alzheimer's disease ($n=10$) and from matched elderly controls ($n=10$). Not only are the ventricles larger in the patients but a prominent ventricular asymmetry (left larger than right) is found in both groups, a feature that only emerges after surface averaging. Specialized approaches for averaging cortical anatomy can also be used to generate population-based maps of brain asymmetry (figure 12d) and to

investigate its alteration in disease (Thompson *et al.* 2000b; Narr *et al.* 2000). Cortical anatomy can also be compared across subjects and its variability encoded to guide the detection of abnormal anatomy (Thompson *et al.* 1997). Figure 12e shows an individual's cortex (brown mesh) overlaid on an average cortical model for a group. Differences in cortical patterns can be encoded by computing a three-dimensional elastic deformation that reconfigures the average cortex into the shape of the individual, matching elements of the gyral pattern exactly (figure 12f). These deformation fields store detailed information on individual deviations and can be averaged across subjects to create three-dimensional variability maps, revealing fundamental patterns of anatomical variability in the brain (figure 12g). The resulting confidence limits on the locations of cortical structures can be used in Bayesian approaches to guide the automated labelling of gyri and sulci (Pitiot *et al.* 2001), and to map profiles of abnormal anatomy in an individual patient or group of subjects (Thompson *et al.* 1995, 1997, 2000b,c; Cao & Worsley 1999).

This strategy has been used to develop atlases and analysis methods for disease states. Specifically, it is both practical and desirable to build disease-specific atlases in order to observe the natural history of a disorder, compare it with normal, age-matched subjects, and to use the disease-specific atlas as a comparison with populations of subjects undergoing conventional or experimental therapies. In this fashion, it is possible to have quantifiable, objective and automated means by which to examine brain structure and function in the normal state and under pathological conditions, as well as during interventions designed to ameliorate or reduce the impact of the disorder. Such an approach using imaging as a surrogate marker of disease

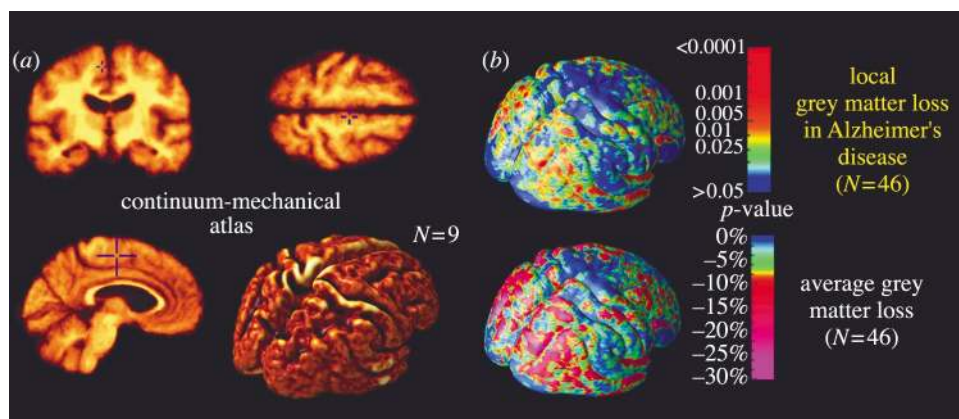


Figure 13. Disease-specific brain atlases. Disease-specific brain atlases reflect the unique anatomy and physiology of a clinical population, in this case an Alzheimer's disease population. Using mathematical strategies to average cortical anatomy across subjects (Thompson *et al.* 2000*b,c*), an average MRI template can be generated for a specific patient group, in this case nine patients with mild to moderate Alzheimer's disease. The cortical pattern indicates clear sulcal widening and atrophic change, especially in temporoparietal cortices. By averaging a measure of grey matter across corresponding cortical regions, the average profiles of grey matter loss can also be mapped. The net reduction in grey matter, in a large patient population relative to controls (Thompson *et al.* 2000*c*) can then be plotted as a statistical map in the atlas. This type of analysis uncovers profiles of early anatomical change in disease. By encoding variations in gyral patterns and grey matter distribution, algorithms can detect a region of abnormal atrophy in the frontal cortex of a dementia patient.

burden may greatly facilitate clinical therapeutic trials by providing objectivity and a quantifiable surrogate endpoint, both of which should increase the cost-effectiveness of these expensive undertakings.

Alzheimer's disease

Probabilistic atlases based on diseased populations (Thompson *et al.* 2000*b,c*) show enormous promise in advancing our understanding of disease. As imaging studies expand into ever-larger patient populations, population-based brain atlases offer a powerful framework to synthesize the results of disparate imaging studies. Disease-specific atlases, for example, are a type of probabilistic atlas specialized to represent a particular clinical group (see Thompson *et al.* 2000*b*, for a review). A disease-specific atlas of brain in Alzheimer's disease has recently been generated to reflect the unique anatomy and physiology of this subpopulation (Thompson *et al.* 1997, 1998, 2000*b*; Mega *et al.* 1997, 1998, 1999). Based on well-characterized patient groups, this atlas contains thousands of structure models as well as composite maps, average templates and visualizations of structural variability, asymmetry and group-specific differences. It also correlates the structural, metabolic, molecular and histological hallmarks of the disease (Mega *et al.* 1997, 1999, 2000). Additional algorithms use information stored in the atlas to recognize anomalies and label structures in new patients. Because they retain information on group anatomical variability, the resulting atlases can identify patterns of altered structure or function and can guide algorithms for knowledge-based image analysis, automated image labelling (Collins *et al.* 1994; Pitiot *et al.* 2001), tissue classification (Zijdenbos & Dawant 1994) and functional image analysis (Dinov *et al.* 2000). At the core of the atlas is an average MRI dataset based on a population of subjects with early dementia (figure 13). Using specialized mathematical approaches for averaging cortical anatomy, the resulting average MRI template has a well-resolved cortical pattern (figure 13*a*) with the mean geometry of the patient group (Thompson *et al.* 2000*b*). Surfaces for the cortex can include the external hull, the grey-white matter

interface or an average of the full cortical thickness. Figure 13*b* represents the external hull.

An example application of this type of atlas is in resolving the average profile of early grey matter loss in an Alzheimer's disease population. It would be ideal, for example, to calibrate the profile of grey matter loss in an individual patient against a normative reference population, for early diagnosis or for clinical trials. Since individual variations in cortical patterning complicate the comparison of grey matter profiles across subjects, an elastic matching technique can be used (driven by 84 structures per brain) that elastically deforms each brain into the group mean geometrical configuration (figure 13*b*). By averaging a measure of grey matter across corresponding regions of cortex, these shape differences are factored out. The net reduction in grey matter, in a large patient population relative to controls ($n = 46$), can then be plotted as a statistical map in the atlas (figure 13*b*; Thompson *et al.* 2000*b*). This type of analysis uncovers important systematic trends, with an early profile of severe grey matter loss detected in temporoparietal cortices, consistent with the early distribution of neuronal loss, metabolic change and perfusion deficits at this stage of Alzheimer's disease. Finally, this local encoding of information on cortical variation can also be exploited to map abnormal atrophy in an individual patient (Thompson *et al.* 1997). Figures 12*f* and 13*b* illustrate the use of a probabilistic atlas to identify a region of abnormal atrophy in the frontal cortex of a dementia patient. Severe abnormality is detected, with a colour code used to indicate the significance of the abnormality (red colours). As expected, corresponding regions in a matched elderly control subject are signalled as normal (figure 13*b*).

(ii) *Cortical surface analysis algorithms*

Cortical surface segmentation

Multiple surface deformation (MSD) is a fully-automated procedure for fitting and unfolding the entire human cortex, using an algorithm which automatically fits a three-dimensional mesh model to the cortical surface extracted from MRI. MSD uses an iterative minimization of a cost function that balances

the distance of the deforming surface from: (i) the target surface; and (ii) the previous iteration surface. Specification of the relative weight of these competing forces allows MSD to range from unconstrained (data driven) deformation to tightly constrained (model preserving) deformation. Further shape preserving constraints are also employed. The initial mesh surface can be chosen arbitrarily to be a simple geometric object, such as a sphere, an ellipsoid or two independently fitted hemispheres (MacDonald *et al.* 1994). Recently, MSD has been extended to allow simultaneous extraction of both inner and outer surfaces of the cortical mantle, using linked concentric mesh models (MacDonald *et al.* 2000). Corresponding vertices in each surface are elastically linked using distance range constraints. Inter-surface cross-intersection and intra-surface self-intersection constraints prevent impossible topologies. These two factors allow for a deeper penetration of the deforming surfaces into the cortical sulci since areas where infolding of the outer (grey–CSF) boundary is indistinct due to partial volume effects are areas where the inner (grey–white) boundary is usually well-distinguished. MSD can operate upon raw image intensity or upon fuzzy-classified tissue maps. Extraction of both surfaces yields a measurement of cortical thickness at each surface vertex. The thickness measurement can be defined in a variety of ways: (i) distance between corresponding vertices; (ii) closest approach of one surface to each vertex of the other surface; (iii) distance between surfaces along the surface normal at each vertex of one surface. These definitions give rise to different absolute values for cortical thickness (closest approach must yield the smallest value, by definition) but the variation in thickness over the whole cortex is generally very similar among the distance measures (MacDonald *et al.* 2000). The method has been applied to a set of 102 MRI volumes from the ICBM database that have been previously mapped automatically into stereotaxic space (Collins *et al.* 1994) and used to generate various group results by averaging the three-dimensional location of corresponding vertices across subjects. The average outer cortical surface obtained when simultaneously fitting both surfaces exhibits a dramatic increase in detail compared with that obtained when fitting only the outer surface, a consequence of the deeper penetration into individual sulci. Since the average cortex can be used as the starting point for mesh-modelling of any individual surface, this is likely to lead to faster and more accurate extraction of individual cortical surfaces in future. Moreover, the average cortical surface is used by some groups to constrain electrophysiological inverse solutions (e.g. Harmony *et al.* 1999) and an improved specification of this surface can be expected to improve that process. The cortical thickness maps exhibit the expected variation in cortical thickness, the temporal poles having the thickest cortex (4–6 mm) and the posterior bank of the central sulcus having the thinnest (1.8–2.5 mm).

This approach has been tested against manual estimates for twenty regions (ten per hemisphere) using 40 brain MRI studies. Validity was determined by an anatomist labelling the CSF–grey and grey–white borders of selected gyri and by allowing the algorithm to determine the CSF–grey and grey–white borders for the same region. The distance between the CSF–grey and grey–white tags determined the cortical thickness at that point. The manual and automatic methods were in agreement for all but four out of 20 regions tested. The four regions where the results were statistically different between the two methods were the insula in both hemispheres, the cuneus and the parahippocampus in the right hemisphere. Thus, the automatic algorithm is valid for most of the cortex and

provides a reasonable alternative to manual *in vivo* measurement except in regions where cortex is adjacent to other grey matter structures.

Sulcal extraction and labelling

We have implemented an automated sulcal extraction and labelling algorithm (SEAL) (LeGoualher *et al.* 1999, 2000). At every voxel on the MSD isosurface, SEAL calculates the two principal curvatures: the mean curvature and the Gaussian curvature. Voxels with negative mean curvature, belonging to sulci, are extracted and pruned to obtain a set of superficial sulcal traces. SEAL extracts the buried sulcus with an ‘active ribbon’ that evolves in three dimensions from a superficial trace to the bottom of a sulcus by optimizing an energy function based on: (i) maximizing distance between starting and current trace position (i.e. for increased penetration); (ii) maximizing distance to any other sulcal voxel (i.e. stay within sulcus); and (iii) minimizing distance from the median sulcal locus. To encode the extracted information, we defined a relational graph structure composed of two main features: arcs and vertices. Arcs contain a surface representing the interior of a sulcus. Points on this surface are expressed in stereotaxic coordinates. For each arc, length, depth and orientation are stored, as well as attributes, e.g. hemisphere, lobe, sulcus type, etc. Each vertex stores its three-dimensional location and its connecting arcs. We have written functions to access this data structure that allow a systematic description of the sulci themselves and their interconnections. Sulcal labelling is performed semi-automatically within DISPLAY by tagging a sulcal trace in the three-dimensional graph and selecting from a menu of candidate labels. The menu is restricted to most likely candidates by the use of spatial priors for sulcal distribution. Given these spatial probability anatomical maps, the user is provided with the probability that the selected arc belongs to a particular sulcus.

(e) Database

Several approaches can be used in the creation of databases to accommodate the diversity of datatypes and structures needed to represent brain structure and function adequately in four dimensions. Whereas a map is a collection of information—a representation of our understanding of the brain—a database is designed with more interactions in mind. Its function is to organize and archive data records and provide an efficient and comprehensive query mechanism. Modern digital maps have only begun to incorporate database functionality.

One of the first database brain maps was developed by Bloom *et al.* (Bloom *et al.* 1990). They created an electronic version of atlas delineations from the Paxinos & Watson (1986) neuro-anatomical atlas of the rat brain. These outlines were equated with coordinates and nomenclature so that the user could request information regarding structural groupings and systems. Since the system was based upon a HyperCard (Apple Computer Corp., Cupertino, CA) database, the user could add information to this anatomical framework as an anatomy laboratory organizer. In a similar vein, Swanson (1992) provided a digital version of his anatomical delineations to his atlas of the rat brain. Cortical connectivity in the macaque monkey also has been organized as a database (Felleman & Van Essen 1991). The most sophisticated attempt in the human brain mapping literature is BrainMap (Fox *et al.* 1994; Fox & Lancaster 1994). This database incorporates a true relational database structure intended to encapsulate data from diverse studies of brain

structure and function. The anatomical framework is based upon the Talairach system and the database relates information about the activation task, the methods, the bibliography and other pertinent data. Usually the image data is excluded and only boundary information is retained. We have chosen to include source image data in the database in a manner that supports both visualization and exploratory query. The query modes include both spatial query of the source image data and query based on a reference anatomical coordinate system.

(i) *Four-dimensional*

A four-dimensional database allows for the intuitive referencing of information by time (age) and place in the nervous system. This works well within a species but requires separate atlases for any given species with links between them to be established only when sufficient information is available to identify true anatomical or functional homologues between the two populations. Once established, the attributes associated with each four-dimensional point could similarly be linked between species-specific atlases and their associated probability estimates.

(ii) *Daemon and BrainMap*

An automated coordinate-based system to retrieve brain labels from the 1988 Talairach atlas, called the Talairach daemon (TD), was previously introduced (Talairach & Tournoux 1988; Lancaster *et al.* 1997). The TD system and its three-dimensional database of labels for the 1988 Talairach atlas were tested for labelling of functional activation foci. The TD system labels were compared with author-designated labels of activation coordinates from over 250 published, functional brain-mapping studies and with manual atlas-derived labels from an expert group using a subset of these activation coordinates. Automated labelling by the TD system compared well with authors' labels, with a 70% or greater label match averaged over all locations. Author-label matching improved to greater than 90% within a search range of ± 5 mm for most sites. An adaptive grey matter (GM) range-search utility was evaluated using individual activations from the M1 mouth region (30 subjects, 52 sites). An 87% label match to Brodmann area labels (BA4 and BA6) was achieved within a search range of ± 5 mm. Using the adaptive GM range search, the TD system's overall match with authors' labels (90%) was better than that of an expert group (80%). When used in concert with authors' deeper knowledge of an experiment, the TD system provides consistent and comprehensive labels for brain activation foci. Additional suggested applications of the TD system include interactive labelling, anatomical grouping of activation foci, lesion-deficit analysis and neuroanatomy education.

5. OTHER ISSUES

(a) *Isolated brain regions*

The more difficult problem than working with whole brain three-dimensional datasets, is that of entering microscopic data from brain sites that are analysed on a regional basis (e.g. the study of the isolated hippocampus). Nevertheless, such data can also be incorporated into the probabilistic reference system and atlas. Such a problem will require landmarks to appropriately localize regional data in the global atlas brain.

Consider a series of post-mortem cryomacrotome human brains that are stained with a series of conven-

tional and commonly used neuroanatomical 'landmark' stains (e.g. Nissl, acetylcholinesterase). These sections would be digitized and sampled at a 20 μ m resolution. The resultant datasets would be warped and entered into the probabilistic atlas as an additional feature. Then consider an investigator who studies gamma aminobutyric acid (GABA) receptors in the human hippocampus. This investigator would like to see where the receptors from the hippocampi of a given epileptic patient population fall with regard to other data in the probabilistic reference system. In preparing the tissue, this investigator would process every *n*th section using one of the 'landmark' stains that are part of the probabilistic atlas. The investigator would then digitize the information from both the GABA receptor sections as well as the 'landmark' stained sections. Using alignment, registration and warping tools that are part of the atlas system, the investigator would register the 'landmark' stained sections with the atlas and then use the same mathematical transformations to enter the GABA receptor information into the hippocampal region of the atlas. Once referenced, database queries and visualization of this new data could be performed in the atlas system. A similar approach allows referencing between newly acquired *in vivo* data and stored post-mortem specimens that should aid in relating functional localization with macroscopic and microscopic anatomy (Rademacher *et al.* 1992; Larsson *et al.* 1999; Naito *et al.* 1999, 2000; Bodegård *et al.* 2000a,b).

(b) *EEG/MEG*

The ICBM atlas is based on neuroanatomy. This is the most fundamental language of communication in neuroscience. As such, it allows appropriate reference and localization to any structure in the brain from any signal source. In the development of the reference system, cross-sectional and tomographic data have been the initial datasets. Once established, however, appropriate vehicles for entering non-tomographic data will be developed. For EEG data, for example, systems already exist to localize scalp electrode placement three-dimensionally, either through the use of a paired tomographic image set or by non-tomographic localization methods (Gevins *et al.* 1994).

(c) *Sociology*

Any endeavour to organize information across laboratories, or especially across an entire field, requires attention to the sociology involved (Koslow 2000). Frustration with existing methods must be high enough and the solutions good enough (in terms of practicality, economics and implementation) that it will be adopted. Such a transition is made easier if rigid new standards are not imposed on the structure or organization of data generated in a given laboratory but rather the tools are available to translate such data into the framework and form required for interaction with the database and atlas. This is a strategy we have employed. Perhaps the most important, if not critical, step is the willingness on the part of the community to share data in all its forms (including raw data) to allow for the full implementation of such a system. Such strategies will require participation of traditional final end products of research (e.g. publication in journals) as well as academic recognition for data

provided to such systems. Lastly, it is always important to have a consensus from the community before embarking on the construction of a complex system such as this one. Wide participation, frequent requests for input and distributed testing of products are all helpful in establishing a successful system that is accepted by the community for which it is intended.

6. LIMITATIONS AND DELIVERABLES

(a) *Deliverables*

A project of this size and scope has a very large overhead at the front end. This results in frustrations for the participants as well as for the community that it is intended to serve. Nevertheless, prior to the release of any individual products (e.g. algorithms, datasets) or the atlas itself, sufficient documentation, validation and a critical mass of test subjects must be acquired in order to be confident of the outcome. We have described our strategy for developing a competitive approach to algorithm development that has at least five phases: theory, initial development, alpha testing, beta testing, and general release. A number of such algorithms have already completed this lengthy process. These include: AIR (Woods *et al.* 1992, 1993, 1998) (now distributed to over 1300 laboratories worldwide), N3 (Sled *et al.* 1997, 1998) (now available at <http://www.bic.mni.mcgill.ca/software/N3/>), an MRI environment simulator (available at <http://www.bic.mni.mcgill.ca/brainweb/>). Algorithms of other components of the ICBM analysis pipeline are well on their way through this competitive process and will be released when complete. The same can be said of datasets. *In vivo* MRI studies as well as cryosection datasets and examples can be found at the ICBM website (<http://www.loni.ucla.edu/ICBM/>).

In an attempt to allow the general neuroimaging and neuroscientific communities to have access to some of the more basic data that has been collected thus far, we are in the process of developing digital libraries. These libraries were described above and will contain raw images (with facial features corrupted), intensity normalized images from multiple pulse sequences as well as normalized and 'scalped' multiple pulse sequences for each subject. This will allow investigators to search for selected subpopulations and use the resultant data for normal controls, methodological developments and many other presently unforeseen uses.

(b) *Limitations*

Every project has its limitations. This one is no different. When faced with the opportunity to evaluate 7000 normal individuals, there is a tendency to be all-inclusive and attempt to collect every potential type of information available. At the onset of this study the contributing investigators met and discussed all of the possible datasets that could be collected from a human subject. The list was long and will not be reiterated here. We opted to start with those datasets that would provide structural imaging of the highest resolution in the largest number of subjects for the best price. We felt that in later years, and in subsequent iterations, it might be possible to add other datasets. In fact, this was done with the addition of functional imaging using fMRI, PET and

event-related potentials. Nevertheless, it was not possible to add information about the vasculature from MR angiography, neurotransmitter systems through PET or single photon emission computed tomography (SPECT) ligand studies, cerebral perfusion through perfusion MRI or PET, chemical information about the brain from MR spectroscopy or datasets that describe major white matter tracts in the brain using diffusion tensor imaging or connectivity using combinations of transcranial magnetic stimulation and PET or fMRI. These are all issues that would be extremely important and valuable to add in the future. Those that have been selected and implemented reflect the basic criteria list noted above as well as the realistic constraints associated with finances, subject risk, time burdens and institutional review board (IRB) criteria. In fact, the reason that only 5800 of the 7000 subjects have DNA samples relates to IRB rules in certain countries with regard to the collection and distribution of genetic materials and information about subjects.

We believe that having neuroanatomy as the basis for building the ICBM probabilistic atlas and reference system was the logical and correct starting point. Other factors that can be added as attributes will be a function of practicality, finances and the interests of the field. They will also be dictated by advances and developments in methodologies.

7. CONCLUSIONS

There is no question that the development of systems and tools such as the probabilistic atlas will have a specific and not insignificant cost associated with them. Also true is the fact that increments in neuroscientific research funding have not kept pace with the growth of the field in terms of numbers of investigators or the magnitude of their projects. Is the cost worth the benefits?

A thoughtful response to this question requires, however, an honest appraisal of the ultimate goals of neuroscientific research. If such research is designed to produce the most accurate understanding of normal brain function and diseases that affect it, then tools that will: (i) enhance the accuracy of results; (ii) enhance the comparison of results between subjects and laboratories; (iii) make more rigorous the confirmation or refutation of data; and (iv) guard against its loss, should have a high priority. A system such as a probabilistic atlas for a given species, or potentially across species, provides a means by which to rigorously store, compare and analyse data over time and between laboratories. Such a system currently does not exist. Furthermore, by virtue of data exchange and comparison, integration within the broad field of neuroscience will be enhanced.

One could rephrase the above statement into a question and ask, 'What would it cost not to develop such integrated systems?'. The costs, in our opinion, would be the progressive and there would be a continued reduction in the value of all funds spent on future neuroscientific research because of the progressively unmanageable amounts and types of data that are generated by neuroscientists. Lacking the tools to manage, compare and analyse these datasets will make funds spent for their acquisition of lesser impact than if such data could be preserved and referenced in an ever-evolving and integrated approach.

Clearly, it would be optimal if funding could come from new sources for systems and approaches to integrate data across not only neuroscience but also computer science, informatics and potentially other related fields. In fact, this is already happening. Contributions to the funding of the initial round of the Human Brain Project (Huerta *et al.* 1993) in the USA came from sources that are both traditional and novel for funding neuroscientific research. By having small contributions from many countries and agencies, the burden on any one country or agency is small but the impact for the neuroscience community is significant. An expanded participation by agencies and contributors outside of traditional pathways as well as the potential for generating new appropriations based on interest in this international effort, will, we hope, result in the creation of systems such as the one described in this report as well as others contemplated or funded through the auspices of the Human Brain Project in the USA, without detracting from traditional neuroscientific funding.

The development of a probabilistic atlas and reference system for the human brain is a formidable goal and one that involves participation from many sites around the world and investigators committed to the end product. The creation of a probabilistic atlas of the human brain is not an exercise in library science. It is a series of fundamental, hypothesis-driven experiments in merging mathematical and statistical approaches with morphological and physiological problems posed with regard to the nervous system. It will create new data and insights into the organization of the human nervous system in health and disease, its development and its evolution. When successful, it will provide previously unprecedented tools for organizing, storing and communicating information about the human brain throughout development, maturation, adult life and old age. It will be a natural prelude to studies of patients with cerebral disorders and provide the first mechanism by which phenotype–genotype–behavioural comparisons can be made on a macroscopic and microscopic level. These results will provide the first insights into the structure–function organization of the human brain across all structures and a wide range of ages in large populations. Its design anticipates the continuing evolution in the quality, resolution and magnitude of data generated by existing technologies that are used to map the human brain and even anticipates that many future technologies, unknown today, will be applicable because the entire system is organized using the architecture of the brain as its guiding principle. The result will allow electronic experimentation and hypothesis generation, facilitated communication among investigators and an objective way of assessing new information gleaned either at scientific meetings or through publications. Developing such a system is an open-ended project with constant evolution, improvement and expansion both in the numbers of subjects included and the range of attributes associated with each. The results should be far more than a data structure and organizational system. Rather, the system should provide new insights and new opportunities for neuroscientists to use data from their own laboratories as well as others to make progress in understanding human brain function in health and disease more rapidly, effectively and efficiently.

This work was supported by a grant from the Human Brain Project (P20-MHDA52176), funded by the National Institute of Mental Health, the National Institute for Drug Abuse, the National Cancer Institute and the National Institute for Neurological Disease and Stroke. For generous support, the authors also wish to thank the Brain Mapping Medical Research Organization, the Pierson–Lovelace Foundation, The Ahmanson Foundation, the Tampkin Foundation, the Jennifer Jones Simon Foundation and the Robson Family. Special thanks to Laurie Carr for the preparation of the manuscript and Andrew Lee for preparation of illustrations. The authors would like to thank the faculties and staff of the participating organizations for their dedication and participation in this program. Finally, we thank the subjects who participated in these investigations for their time, interest and commitment.

REFERENCES

- Ahlfors, S. P. (and 10 others) 1999 Spatiotemporal activity of a cortical network for processing visual motion revealed by MEG and fMRI. *J. Neurophysiol.* **82**, 2545–2555.
- Amdimts, K., Malikovic, A., Mohlberg, H., Schormann, T. & Zilles, K. 2000 Brodmann's areas 17 and 18 brought into stereotaxic space—where and how variable? *NeuroImage* **11**, 66–84.
- Amunts, K., Klingberg, T., Binkofski, F., Schormann, T., Seitz, R. J., Roland, P. E. & Zilles, K. 1998 Cytoarchitectonic definition of Broca's region and its role in functions different from speech. *NeuroImage* **7**, 8.
- Amunts, K., Schleicher, A., Biirgel, U., Mohlberg, H., Uylings, H. B. M. & Zilles, K. 1999 Broca's region *y* revisited: cytoarchitecture and intersubject variability. *J. Comp. Neurol.* **412**, 319–341.
- Arnold, J. B. (and 10 others) 2001 Quantitative and qualitative evaluation of six algorithms for correcting intensity non-uniformity effects. *NeuroImage* **13**, 931–943.
- Ashburner, J., Neelin, P., Collins, D. L., Evans, A. & Friston, K. 1997 Incorporating prior knowledge into image registration. *NeuroImage* **6**, 344–352.
- Ayle, C. J., Supek, S., George, J. S., Ranken, D., Lewine, J., Sanders, J., Best, E., Tice, W., Flynn, E. R. & Wood, C. C. 1996 Retinotopic organization of human visual cortex: departures from the classical model. *Cerebral Cortex* **6**, 354–361.
- Bailey, P. & von Bonin, G. 1951 *The isocortex of man*. Urbana University Press.
- Bajcsy, R. & Kovacic, S. 1989 Multi-resolution elastic matching. *Comput. Vision, Graphics, Image Process.* **46**, 1–21.
- Bartley, A. J., Jones, D. W. & Weinberger, D. R. 1997 Genetic variability of human brain size and cortical gyral patterns. *Brain* **120**, 257–269.
- Bloom, F. E. 1996 The multidimensional database and neuroinformatics requirements for molecular and cellular neuroscience. *NeuroImage* **4**, S12–S13.
- Bloom, F. E., Young, W. G. & Kirn, Y. M. 1990 *BrainBrowser: HyperCard application for the Macintosh*. San Diego, CA: Academic Press.
- Bodegård, A., Geyer, S., Naito, E., Zilles, K. & Roland, P. E. 2000a Somatosensory areas in man activated by moving stimuli: cytoarchitectonic mapping and PET. *NeuroReport* **11**, 187–191.
- Bodegård, A., Ledberg, A., Geyer, S., Naito, E., Larsson, J., Zilles, K. & Roland, P. 2000b Object shape differences reflected by somatosensory cortical activation in human. *J. Neurosci.* **20**, 1–5.
- Brodmann, K. 1909 *Vergleichende Lokalisationslehre der Großhirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues*. Leipzig: Barth JA.
- Broit, C. 1981 Optimal registration of deformed images. PhD dissertation, University of Pennsylvania, USA.

- Bro-Nielsen, M. & Gramkow, C. 1996 Fast fluid registration of medical images. In *Proc. Fourth Int. Conf. on Visualization in Biomedical Computing* ((VBC'96) ed. K. H. Hohne & R. Kikinis), Hamburg, Germany, 22–25 September 1996, pp. 267–276. Berlin: Springer-Verlag.
- Cannestra, A. F., Santori, E. M., Holmes, C. J. & Toga, A. W. 1997 A three-dimensional multimodality map of the nemistrina monkey. *Brain Res. Bull.* **5**, 147–153.
- Cao, J. & Worsley, K. J. 1999 The geometry of the Hotelling's T² random field with applications to the detection of shape changes. *Anals Statistics* **27**, 925–942.
- Carman, G. J., Drury, H. A. & Van Essen, D. C. 1995 Computational methods for reconstructing and unfolding the cerebral cortex. *Cerebral Cortex* **5**, 506–517.
- Christensen, G., Rabbitt, R. D., Miller, M. I. 1993 A deformable neuroanatomy textbook based on viscous fluid mechanics. Invited paper. In *Proc. 1993 Conf. Info. Sci. Syst.* (ed. J. Prince & T. Runolfsson), pp. 211–216. Baltimore, MD: Johns Hopkins University Press.
- Christensen, G. E., Rabbitt, R. D., Miller, M. I., Joshi, S. C., Grenander, U., Coogan, T. A. & Van Essen, D. C. 1995 Topological properties of smooth anatomic maps. In *Information processing in medical imaging* (ed. Y. Bizais, C. Barillot & R. DiPaola), June 1995, pp. 101–112. Boston, MA: Kluwer Academic.
- Christensen, G. E., Rabbitt, R. D. & Miller, M. I. 1996 Deformable templates using large deformation kinematics. *IEEE Trans. Image Processing* **5**, 1435–1447.
- Clark, S. & Miklossy, J. 1990 Occipital cortex in man: organization of callosal connections, related myelo- and cytoarchitecture, and putative boundaries of functional visual areas. *J. Comp. Neurol.* **298**, 188–214.
- Collins, D. L., Neelin, P., Peters, T. M., Evans, A. C. 1994 Automatic 3D registration of MR volumetric data in standardized Talairach space. *J. Comp. Assist. Tomogr.* **8**, 192–205.
- Collins, D. L., Holmes, C. J., Peters, T. M. & Evans, A. C. 1995 Automatic 3D model-based neuroanatomical segmentation. *Hum. Brain Mapp.* **3**, 190–208.
- Dale, A. M., Pisch, B., Sereno, M. I. 1999 Cortical surface-based analysis. I. Segmentation and surface reconstruction. *NeuroImage* **9**, 179–194.
- Davatzikos, C. 1996 Spatial normalization of 3D brain images using deformable models. *J. Comp. Assist. Tomogr.* **20**, 656–665.
- Davatzikos, C. 1997 Spatial transformation and registration of brain images using elastically deformable models. *Comput. Vis. Image Und.* **66**, 207–222.
- Dinov, I. D., Mega, M. S., Thompson, P. M., Lee, L., Woods, R. P., Holmes, C. J., Summers, D. W. & Toga, A. W. 2000 Analyzing functional brain images in a probabilistic atlas: a validation of subvolume thresholding. *J. Comp. Assist. Tomogr.* **24**, 128–138.
- Dumoulin, S. O., Bittar, R. G., Kabani, N. J., Baker, C. L., LeGoualher, G., Pike, G. B. & Evans, A. C. 2000 A new neuroanatomical landmark for the reliable identification of human area V5/MT: a quantitative analysis of sulcal patterning. *Cerebral Cortex* **10**, 454–463.
- Duta, N., Sonka, M. & Jain, A. K. 1999 Learning shape models from examples using automatic shape clustering and Procrustes analysis. In *Proc. 16th Int. Conf. on Information Processing in Medical Imaging* (IPMI'99) (ed. A. Kuba, M. Sámal & A. Todd-Pokropek), pp. 370–375. Lecture Notes in Computer Science, vol. 1613. Berlin: Springer-Verlag.
- Evans, A. C., Frank, J. A., Antel, J. & Miller, D. H. 1997 The role of MRI in clinical trials of multiple sclerosis: comparison of image processing techniques. *Anals Neurol.* **41**, 125–132.
- Evans, A. C., Collins, D. L. & Holmes, C. J. 1996 Automatic 3D regional MRI segmentation and statistical probability anatomy maps. In *Quantification of brain function using PET* (ed. R. Myers, V. Cunningham, D. Bailey & T. Jones), pp. 123–130. San Diego, CA: Academic Press.
- Evans, A. C., Dai, W., Collins, D. L., Neelin, P. & Marrett, S. 1991 Warping of a computerized 3D atlas to match brain image volumes for quantitative neuroanatomical and functional analysis. *Proc. SPIE Med. Imag.* **1445**, 236–247.
- Felleman, D. & Van Essen, D. 1991 Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex* **1**, 1–47.
- Filimonoff, I. N. 1932 Über die Variabilität der Großhirnrindenstruktur. Mitteilung II—Regio occipitalis beim erwachsenen Menschen. *J. Psychol. Neurol.* **44**, 2–96.
- Fischl, B., Sereno, M. I., Tootell, R. B., Dale, A. M. 1999 High-resolution intersubject averaging and a coordinate system for the cortical surface. *Hum. Brain Mapp.* **8**, 272–284.
- Flechsig, P. 1920 *Anatomie des menschlichen Gehirns und Rückenmarks*. Leipzig: Thieme.
- Fox, P. T. & Mintun, M. A. 1989 Noninvasive functional brain mapping by change-distribution analysis of averaged PET images of H₂¹⁵O tissue activity. *J. Neurosci.* **7**, 913–922.
- Fox, P. T. & Lancaster, J. L. 1994 Neuroscience on the net. *Science* **266**, 994–996.
- Fox, P. T., Mikiten, S., Davis, G. & Lancaster, J. 1994 Brain map: a database of human functional brain mapping. In *Functional neuroimaging* (ed. R. W. Thatcher, M. Hallett, T. Zeffiro, E. R. John & M. Huerta), pp. 95–105. San Diego, CA: Academic Press.
- Friston, K. J., Frith, C. D., Liddle, P. F. & Frackowiak, R. S. J. 1991 Comparing functional (PET) images: the assessment of significant change. *J. Cerebr. Blood Flow Metab.* **11**, 690–699.
- Friston, K. J., Holmes, A. P., Worsley, K. J., Poline, J. P., Frith, C. D. & Frackowiak, R. S. J. 1995. Statistical parametric maps in functional imaging: a general linear approach. *Hum. Brain Mapp.* **2**, 189–210.
- Galaburda, A. M. 1980 La región de Broca: observaciones anatómicas hechas un siècle après la mort de son découvreur. *Revue Neurologique* **136**, 609–616.
- Galaburda, A. M., LeMay, M., Kemper, T. L. & Geschwind, N. 1978 Right–left asymmetries in the brain. *Science* **199**, 852–856.
- Gee, J. C., Reivich, M. & Bajcsy, R. 1993 Elastically deforming an atlas to match anatomical brain images. *J. Comp. Assist. Tomogr.* **17**, 225–236.
- Gee, J. C., LeBriquer, L., Barillot, C., Haynor, D. R. & Bajcsy, R. 1995 Bayesian approach to the brain–image matching problem. *Proc. SPIE Medical Imag.* **2434**, 145–156.
- Geschwind, N. & Levitsky, W. 1968 Human brain: left–right asymmetries in temporal speech region. *Science* **161**, 186–187.
- Gevens, A. S., Lee, J., Martin, N., Reutter, R., Desmond, J. & Brickett, P. 1994 High resolution EEG: 124-channel recording, spatial deblurring and MRI integration methods. *Electroen Clin. Neuro.* **90**, 337–358.
- Geyer, S., Ledberg, A., Schleicher, A., Kinomura, S., Schormann, T., Bürgel, U., Klingberg, T., Larsson, J., Zilles, K. & Roland, P. E. 1996 Two different areas within the primary motor cortex of man. *Nature* **382**, 805–807.
- Geyer, S., Schleicher, A., Zilles, K. 1997 The somatosensory cortex of human: cytoarchitecture and regional distributions of receptor-binding sites. *NeuroImage* **6**, 27–45.
- Geyer, S., Schleicher, A. & Zilles, K. 1999 Areas 3a, 3b, and 1 of human primary somatosensory cortex. 1. Microstructural organization and interindividual variability. *NeuroImage* **10**, 63–83.
- Geyer, S., Schormann, T., Mohlberg, H. & Zilles, K. 2000 Areas 3a, 3b, and 1 of human primary somatosensory cortex. 2. Spatial normalization to standard anatomical space. *NeuroImage* **11**, 684–696.
- Gur, R. C., Packer, I. K., Hungerbuhler, J. P., Reivich, M., Obrist, W., Amamek, W. & Sackeim, H. 1980 Differences in the distribution of gray and white matter in human cerebral hemispheres. *Science* **207**, 1226–1228.

- Harmony, T., Fernandez, T., Silva, J., Bosch, J., Valdes, P., Fernandez-Bouzas, A., Galan, L., Aubert, E. & Rodriguez, D. 1999 Do specific EEG frequencies indicate different processes during mental calculation? *Neurosci. Lett.* **266**, 25–28.
- Hayes, T. L. & Lewis, D. A. 1995 Anatomical specialization of the anterior motor speech area: hemispheric differences in magnopyramidal neurons. *Brain Lang.* **49**, 289–308.
- Hayes, T. L. & Lewis, D. A. 1996 Magnopyramidal neurons in the anterior motor speech region. *Arch. Neurol.* **53**, 1277–1283.
- Holmes, C. J., Hoge, R., Collins, L., Woods, R., Toga, A. W. & Evans, A. C. 1998 Enhancement of magnetic resonance images using registration for signal averaging. *J. Comp. Assist. Tomogr.* **22**, 324–333.
- Huerta, M., Koslow, S. & Leshner, A. 1993 The Human Brain Project: an international resource. *TINS* **16**, 436–438.
- Ingvar, M., Eriksson, L., Greitz, T., Stone-Elander, S., Dahlbom, M., Rosenqvist, G., af Trampe, P. & von Euler, C. 1994 Methodological aspects of brain activation studies: cerebral blood flow determined with [15O]butanol and positron emission tomography. *J. Cerebr. Blood Flow Metab.* **14**, 628–38.
- Jacobs, B., Batal, H. A., Lynch, B., Ojemann, G., Ojemann, L. M. & Scheibel, A. B. 1993 Quantitative dendritic and spine analysis of speech cortices: a case study. *Brain Lang.* **44**, 239–253.
- Kamber, M., Shinghal, R., Collins, D. L., Francis, G. S. & Evans, A. C. 1995 Model-based 3D segmentation of multiple sclerosis lesions in magnetic resonance brain images. *IEEE Trans. Med. Imaging* **14**, 442–453.
- Kamber, M., Collins, D. L., Shinghal, R., Francis, G. S. & Evans, A. C. 1992 Model-based 3D segmentation of multiple sclerosis lesions in dual-echo MRI data. *Proc. SPIE Vis. Biomed. Comp.* **1808**, 590–600.
- Koslow, S. H. 2000 Should the neuroscience community make a paradigm shift to sharing primary data? *Nature Neurosci.* **3**, 863–865.
- Lancaster, J. L., Rainey, L. H., Summerlin, J. L., Freitas, C. S., Fox, P. T., Evans, A. C., Toga, A. W. & Mazziotta, J. C. 1997 Automated labeling of the human brain: a preliminary report on the development and evaluation of a forward-transform method. *Hum. Brain Mapp.* **5**, 238–242.
- Larsson, J., Amunts, K., Gulyás, B., Malikovic, A., Zilles, K. & Roland, P. E. 1999 Neuronal correlates of real and illusory contour perception: functional anatomy with PET. *Euro. J. Neurosci.* **11**, 4024–4036.
- LeGoualher, G., Procyk, E., Collins, D. L., Venugopal, R., Barillot, C. & Evans, A. C. 1999 Automated extraction and variability analysis of sulcal neuroanatomy. *IEEE Trans. Med. Imaging* **18**, 206–217.
- LeGoualher, G., Argenti, A., Duyme, M., Baare, W., Hulshoff Pol, H., Barillot, C. & Evans, A. C. 2000 Statistical sulcal shape comparisons: application to the detection of genetic encoding of the central sulcus shape. *NeuroImage* **11**, 564–574.
- Leonard, C. M., Kuldau, J. M., Breier, J. I., Zuffante, P. A., Gautier, E. R., Heron, D.-C., Lavery, E. M., Packing, J., Williams, S. A. & DeBose, C. A. 1999 Cumulative effect of anatomical risk factors for schizophrenia: an MRI study. *Biol. Psychiatry* **46**, 374–382.
- Lohmann, G. 1998 Extracting line representations of sulcal and gyral patterns in MR images of the human brain. *IEEE Trans. Med. Imaging* **17**, 1040–1048.
- MacDonald, D. 1998 A method for identifying geometrically simple surfaces from three-dimensional images. PhD dissertation, McGill University, USA.
- MacDonald, D., Avis, D. & Evans, A. C. 1994 Multiple surface identification and matching in magnetic resonance imaging. *Proc. SPIE Vis. Biomed. Comp.* **2359**, 160–169.
- MacDonald, D., Kabani, N., Avis, D. & Evans, A. C. 2000 Automated 3D extraction of inner and outer surfaces of cerebral cortex from MRI. *NeuroImage* **12**, 340–356.
- Mazziotta, J. C. 1984 Physiological neuroanatomy: functional brain imaging presents a new problem to an old discipline. *J. Cerebr. Blood Flow Metab.* **4**, 481–484.
- Mazziotta, J. C., Huang, S. C., Phelps, M. E., Carson, R. E., MacDonald, N. S. & Mahoney, K. 1985 A non-invasive positron computed tomography technique using oxygen-15 labeled water for the evaluation of neurobehavioral task batteries. *J. Cerebr. Blood Flow Metab.* **5**, 70–78.
- Mazziotta, J. C., Toga, A. W., Evans, A. C., Fox, P. & Lancaster, J. 1995 A probabilistic atlas of the human brain: theory and rationale for its development. *NeuroImage* **2**, 89–101.
- Mazziotta, J. C., Toga, A. W., Evans, A., Fox, P. & Lancaster, J. 1995a Digital brain atlases. *Trends Neurosci.* **18**, 210–211.
- Mega, M. S., Chen, S., Thompson, P. M., Woods, R. P., Karaca, T. J., Tiwari, A., Vinters, H., Small, G. W. & Toga, A. W. 1997 Mapping pathology to metabolism: coregistration of stained whole brain sections to PET in Alzheimer's disease. *NeuroImage* **5**, 147–153.
- Mega, M. S., Thompson, P. M., Cummings, J. L., Back, C. L., Xu, M. L., Zohoori, S., Goldkom, A., Moussai, J., Fairbanks, L., Small, G. W. & Toga, A. W. 1998 Sulcal variability in the Alzheimer's brain: correlations with cognition. *Neurology* **50**, 145–151.
- Mega, M. S., Chu, T., Mazziotta, J. C., Trivedi, K. H., Thompson, P. M., Shah, A., Cole, G., Frautschy, S. A. & Toga, A. W. 1999 Mapping biochemistry to metabolism: FDG-PET and amyloid burden in Alzheimer's disease. *NeuroReport* **10**, 2911–2917.
- Mega, M. S., Lee, L., Dinov, I. D., Mishkin, F., Toga, A. W., Cummings, J. L. 2000 Cerebral correlates of psychotic symptoms in Alzheimer's disease. *J. Neurol. Neurosurg. Psychiatry* **69**, 167–171.
- Merker, B. 1983 Silver staining of cell bodies by means of physical development. *J. Neurosci.* **9**, 235–241.
- Miller, M. I., Christensen, G. E., Amit, Y. & Grenander, U. 1993 Mathematical textbook of deformable neuroanatomies. *Proc. Natl Acad. Sci. USA* **90**, 11944–11948.
- Naito, E., Ehrsson, H. H., Geyer, S., Zilles, K. & Roland, P. E. 1999 Illusory arm movements activate cortical motor areas: a positron emission tomography study. *J. Neurosci.* **19**, 6134–6144.
- Naito, E., Kinomura, S., Geyer, S., Kawashima, R., Roland, P. E. & Zilles, K. 2000 Fast reaction to different sensory modalities activates common fields in the motor areas, but the anterior cingulate cortex is involved in the speed of reaction. *J. Neurophysiol.* **83**, 1701–1709.
- Narr, K. L., Thompson, P. M., Sharma, T., Moussai, J., Cannestra, A. F. & Toga, A. W. 2000 Mapping morphology of the corpus callosum in schizophrenia. *Cerebral Cortex* **10**, 40–49.
- Neelin, P. D., MacDonald, D., Collins, D. L. & Evans, A. C. 1998 The MINC file format: from bytes to brains. *NeuroImage* **7**, S786.
- Ono, M., Kubik, S. & Abemathy, C. 1990 *Atlas of the cerebral sulci*. New York: Thieme Medical Publishers, Stuttgart, Germany.
- Paus, T., Otaky, N., Caramanos, Z., MacDonald, D., Zijdenbos, A., D'Avirro, D., Gutmans, D., Holmes, C. J., Tomaiuolo, F. & Evans, A. C. 1996a *In vivo* morphometry of the intrasulcal gray matter in the human cingulate, paracingulate and superior-rostral sulci: hemispheric asymmetries, gender differences, and probability maps. *J. Comp. Neurol.* **376**, 664–673.
- Paus, T., Tomaiuolo, F., Otaky, N., MacDonald, D., Petrides, M., Atlas, J., Morris, R. & Evans, A. C. 1996b Human

- cingulate and paracingulate sulci: Pattern, variability, asymmetry, and probabilistic map. *Cerebral Cortex* **6**, 207–214.
- Paxinos, G. & Watson, C. 1986 *The rat brain in stereotaxic coordinates*. Sydney: Academic Press.
- Payne, B. A. & Toga, A. W. 1990 Surface mapping brain function on 3D models. *IEEE Comp. Graph.* **10**, 33–41.
- Penhune, V. B., Zatorre, R. J., MacDonald, J. D. & Evans, A. C. 1996 Interhemispheric anatomical differences in human primary auditory cortex: Probabilistic mapping and volume measurement from MR scans. *Cerebral Cortex* **6**, 617–672.
- Pennisi, E. 1999 Keeping genome databases clean and up to date. *Science* **286**, 447–450.
- Pitiot, A., Toga, A. W. & Thompson, P. M. 2001 Spatially and temporally adaptive elastic template matching. *IEEE Trans. Med. Imaging*. (Submitted.)
- Polyak, S. L. 1957 *The vertebrate visual system*. University of Chicago Press.
- Rabbitt, R. D., Weiss, J. A., Christensen, G. E. & Miller, M. I. 1995 Mapping of hyperelastic deformable templates using the finite element method. *Proc. SPIE* **2573**, 252–265.
- Rademacher, J., Galaburda, A., Kennedy, D., Filipek, P. & Caviness, V. 1992 Human cerebral cortex: localization, parcellation, and morphometry with magnetic resonance imaging. *J. Cogn. Neurosci.* **4**, 352–374.
- Rademacher, J., Caviness, V. S., Steinmetz, H., Galaburda, A. M. 1993 Topographical variation of the human primary cortices: implications for neuroimaging, brain mapping and neurobiology. *Cerebr. Cortex* **3**, 313–329.
- Rizzo, G., Gilardi, M. C., Prinster, A., Grassi, F., Scotti, G., Cerutti, S. & Fazio, F. 1995 An elastic computerized brain atlas for the analysis of clinical PET/SPET data. *Eur. J. Nucl. Med.* **22**, 1313–1318.
- Roland, P. E. & Zilles, K. 1994 Brain atlases: a new research tool. *Trends Neurosci.* **17**, 458–467.
- Roland, P. E. & Zilles, K. 1996 The developing European computerized human brain database for all imaging modalities. *NeuroImage* **4**, 39–47.
- Roland, P. E. & Zilles, K. 1998 Structural divisions and functional fields in the human cerebral cortex. *Brain Res. Rev.* **26**, 87–105.
- Sandor, S. & Leahy, R. 1994 Matching deformable atlas models to pre-processed magnetic resonance brain images. *Proc. IEEE Image Process.* **3**, 686–690.
- Sandor, S. R. & Leahy, R. M. 1995 Towards automated labeling of the cerebral cortex using a deformable atlas. In *Proc 14th Int. Conf. Inform. Process. Me. Imaging* (ed. Y. Bizais, C. Barillot & R. Di Paola), pp. 127–138. Boston, MA: Kluwer Academic Press.
- Sandor, S. & Leahy, R. 1997 Surface-based labeling of cortical anatomy using a deformable atlas. *IEEE Trans. Med. Imaging* **16**, 41–54.
- Scheibel, A. B., Paul, L. A., Fried, I., Forsythe, A. B., Tomiyasu, U., Wechsler, A., Kao, A. & Slotnick, J. 1985 Dendritic organization of the anterior speech area. *Exp. Neurol* **87**, 109–117.
- Schleicher, A., Amunts, K., Geyer, S., Morosan, P. & Zilles, K. 1999 Observer-independent method for microstructural parcellation of cerebral cortex: a quantitative approach to cytoarchitectonics. *NeuroImage* **9**, 165–177.
- Schmidt, D. M., George, J. S. & Wood, C. C. 1999 Bayesian inference applied to the electromagnetic inverse problem. *Hum. Brain Mapp.* **7**, 195–212.
- Schormann, T. & Zilles, K. 1997 Limitations of the principal axes theory. *IEEE Trans. Med. Imaging* **16**, 942–947.
- Schormann, T. & Zilles, K. 1998 Three-dimensional linear and nonlinear transformations: an integration of light microscopical and MRI data. *Hum. Brain Mapp.* **6**, 339–347.
- Schormann, T., Dabringhaus, A. & Zilles, K. 1995 Statistics of deformations in histology and improved alignment with MRI. *IEEE Trans. Med. Imaging* **14**, 25–35.
- Schormann, T., Henn, S. & Zilles, K. 1996 A new approach to fast elastic alignment with application to human brains. *Proc. Vis. Biomed. Comp.* **4**, 337–342.
- Schormann, T., Dabringhaus, A. & Zilles, K. 1997 Extension of the principal axis theory for the determination of affine transformations. In *Informatik aktuell.*, 384–391. Springer.
- Seitz, R. J., Bohm, C., Greitz, T., Roland, P. E., Eriksson, L., Blomqvist, G., Rosenqvist, G. & Nordell, B. 1990 Accuracy and precision of the computerized brain atlas programme for localization and quantification in positron emission tomography. *J. Cerebr. Blood Flow Metab.* **10**, 443–457.
- Simonds, R. J. & Scheibel, A. B. 1989 The postnatal development of the motor speech area: a preliminary study. *Brain Lang.* **37**, 43–58.
- Sled, J., Zijdenbos, A. & Evans, A. 1997 A comparison of retrospective intensity non-uniformity correction methods for MRI. In *Proc. 15th Int. Conf. on Information Processing and Medical Imaging (IPMI'97), June 1997, Poultney, VT*, pp. 459–464. Berlin: Springer-Verlag.
- Sled, J., Zijdenbos, A. & Evans, A. 1998 A non-parametric method for automatic correction of intensity non-uniformity in MRI data. *IEEE Trans. Med. Imaging* **17**, 87–97.
- Sowell, E. R., Mattson, S. N., Thompson, P. M., Jernigan, T. L., Riley, E. P. & Toga, A. W. 2001 Mapping callosal morphology and cognitive correlates: effects of heavy prenatal alcohol exposure. *Neurology* **57**. (In the press.)
- Spitzer, V. M. & Whitlock, D. G. 1992 High resolution imaging of the human body. *J. Biol. Photogr.* **60**, 167–172.
- Steinmetz, H., Volkman, J., Jancke, L. & Freund, H. 1991 Anatomical left–right asymmetry of language-related temporal cortex is different in left and right handers. *Annls Neurol.* **29**, 315–319.
- Strother, S. C., Anderson, J. R., Xu, X.-L., Low, J.-S., Boar, D. C. & Rottenberg, D. A. 1994 Quantitative comparisons of image registration techniques based on high-resolution MRI of the brain. *J. Comp. Assist. Tomogr.* **18**, 954–962.
- Swanson, L. W. 1992 *Computer graphics file. Version 1.0*. Amsterdam: Elsevier-Science.
- Talairach, J. & Tournoux, P. 1988 Principe et technique des études anatomiques. In *Co-planar stereotaxic atlas of the human brain—3-dimensional proportional system: an approach to cerebral imaging*. New York: Thieme Medical Publishers.
- Thirion, J.-P. 1995 Fast non-rigid matching of medical images. Rapport de Recherche no. 2547. Institut National de Recherche en Informatique et en Automatique, Le Chesney, France.
- Thompson, P. M. & Toga, A. W. 1996 A surface-based technique for warping 3-dimensional images of the brain. *IEEE Trans. Med. Imaging* **15**, 1–16.
- Thompson, P. M. & Toga, A. W. 1997 Detection, visualization and animation of abnormal anatomic structure with a deformable probabilistic brain atlas based on random vector field transformations. *Med. Image Anal.* **1**, 271–294.
- Thompson, P. M. & Toga, A. W. 1998 Anatomically driven strategies for high-dimensional brain image warping and pathology detection. In *Brain warping* (ed. A. W. Toga), pp. 311–336. San Diego, CA: Academic Press.
- Thompson, P. M. & Toga, A. W. 2000 Warping strategies for inter-subject registration. In *Handbook of medical image processing* (ed. Isaac Bankman), pp. 569–601. San Diego, CA: Academic Press.
- Thompson, P. M., Schwartz, C., Lin, R. T., Khan, A. A., Toga, A. W. & Collins, R. C. 1995 3D statistical analysis of sulcal variability in the human brain using high-resolution cryosection images. *Proc. Soc. Neurosci.* **21**, 154.
- Thompson, P. M., Schwartz, C. & Toga, A. W. 1996a High-resolution random mesh algorithms for creating a probabilistic 3D surface atlas of the human brain. *NeuroImage* **3**, 19–34.
- Thompson, P., Schwartz, C., Lin, R. T., Khan, A. A., Toga, A. W. 1996b 3D statistical analysis of sulcal variability in the human brain. *J. Neurosci.* **16**, 4261–4274.

- Thompson, P. M., MacDonald, D., Mega, M. S., Holmes, C. J., Evans, A. C. & Toga, A. W. 1997 Detection and mapping of abnormal brain structure with a probabilistic atlas of cortical surfaces. *J. Comp. Assist. Tomogr.* **21**, 567–581.
- Thompson, P. M., Moussai, J., Khan, A. A., Zohoori, S., Goldkom, A., Mega, M. S., Small, G. W., Cummings, J. L. & Toga, A. W. 1998 Cortical variability and asymmetry in normal aging and Alzheimer's disease. *Cerebral Cortex* **8**, 492–509.
- Thompson, P. M., Giedd, J. N., Woods, R. P., MacDonald, D., Evans, A. C. & Toga, A. W. 2000a Growth patterns in the developing brain detected by using continuum mechanical tensor maps. *Nature* **404**, 190–193.
- Thompson, P. M., Mega, M. S. & Toga, A. W. 2000b Disease-specific brain atlases. In *Brain mapping: the disorders* (ed. A. W. Toga, J. C. Mazziotta & R. S. J. Frackowiak), pp. 131–177. San Diego, CA: Academic Press.
- Thompson, P. M., Wood, R. P., Mega, M. S. & Toga, A. W. 2000c Mathematical/computational challenges in creating deformable and probabilistic atlases of the human brain. *Hum. Brain Mapp.* **9**, 81–92.
- Thurfjell, L., Bohm, C., Greitz, T. & Eriksson, L. 1993 Transformations and algorithms in a computerized brain atlas. *IEEE Trans. Nucl. Sci.* **40**, 1167–1191.
- Toga, A. W., Thompson, P. M., Holmes, C. J. & Payne, B. A. 1996 Informatics and Computational Neuroanatomy. *J. Am. Med. Informatics Assn* **6**, 299–303.
- Toga, A. W., Goldkom, A., Ambach, K., Chao, K., Quinn, B. C. & Yao, P. 1997 Postmortem cryosectioning as an anatomic reference for human brain mapping. *Comp. Med. Image Graphics* **21**, 131–141.
- Toga, A. W. & Amicar-Sulze, T. L. 1987 Digital image reconstruction for the study of brain structure and function. *J. Neurosci. Meth.* **20**, 7–21.
- Ungerleider, L. G. & Desimone, R. 1986 Cortical connections of visual area MT in the macaque. *J. Comp. Neurol.* **248**, 190–222.
- Van Essen, D. C. & Drury, H. A. 1997 Structural and functional analyses of human cerebral cortex using a surface-based atlas. *J. Neurosci.* **17**, 7079–7102.
- Van Essen, D. C. & Drury, H. A., Joshi, S. & Miller, M. I. 1998 Functional and structural mapping of human cerebral cortex: solutions are in the surfaces. *Proc Natl Acad. Sci. USA* **95**, 788–95.
- Viola, P. A. & Wells, W. M. 1997 Alignment by maximization of mutual information. *Int. J. Comp. Vis.*, **24**, 137–154.
- Wang, J. Z., Wiederhold, G. & Firschein, O. 1997 System for screening objectionable images using Daubechies' wavelets and color histograms, interactive distributed multimedia systems and telecommunication services. *Proc. 4th Euro Workshop (IDMS'97)* (ed. R. Steinmetz & L. C. Wolf). Darmstadt, Germany: Springer-Verlag.
- Watson, J. D., Myers, R., Frackowiak, R. S., Hajnal, J. V., Mazziotta, J. C., Shipp, S. & Zeki, S. 1993 Area V5 of the human brain from a combined study using positron emission tomography and magnetic resonance imaging. *Cerebral Cortex* **3**, 79–94.
- Wells, W. M., Viola, P., Atsumi, H., Nakajima, S. & Kikinis, R. 1997 Multi-modal volume registration by maximization of mutual information. *Med. Image Anal.* **1**, 35–51.
- Wertheim, S. L. 1989 The brain database: a multimedia neuroscience database for research and teaching. In *Proc. 13th Ann. Symp. Comp. Appl. Med. Care* (ed. L. C. Kingsland III). Los Alamitos, CA: IEEE Computer Society Press.
- Woods, R. P., Grafton, S. T., Holmes, C. J., Cherry, S. R. & Mazziotta, J. C. 1998 Automated image registration. I. General methods and intrasubject, intramodality validation. *J. Comput. Assist. Tomogr.* **22**, 139–152.
- Woods, R. P., Cherry, S. R. & Mazziotta, J. C. 1992 A rapid automated algorithm for accurately aligning and reslicing positron emission tomography images. *J. Comp. Assist. Tomogr.* **16**, 620–633.
- Woods, R. P., Mazziotta, J. C. & Cherry, S. R. 1993 MRI–PET registration with automated algorithm. *J. Comp. Assist. Tomogr.* **17**, 536–546.
- Zeki, S., Watson, J. D. G., Lueck, C. J., Friston, K. J., Kennard, C. & Frackowiak, R. S. J. 1991 A direct demonstration of functional specialization in human visual cortex. *J. Neurosci.* **11**, 641–649.
- Zhou, Y. & Toga, A. W. 1999 Efficient skeletonization of volumetric objects. *IEEE Trans. Vis. Computer Graphics* **5**, 196–209.
- Zhou, Y., Kaufman, A. & Toga, A. W. 1998 Three-dimensional skeleton and centerline generation based on an approximate minimum distance field. *Visual Computer* **14**, 303–314.
- Zihl, J., von Cramon, D., Mai, D. & Schmid, C. 1991 Disturbance of movement vision after bilateral posterior brain damage. Further evidence and follow up observations. *Brain* **114**, 2235–2252.
- Zijdenbos, A. P. & Dawant, B. M. 1994 Brain segmentation and white matter lesion detection in MR images. *Crit. Rev. Biomed. Engng* **22**, 401–465.
- Zijdenbos, A. P., Evans, A. C., Riahi, F., Sled, J. G., Chui, H.-C., Kollokian V. 1996 Automatic quantification of multiple sclerosis lesion volume using stereotaxic space. In *Proc. 4th Int. Conf. Visualization Biomed. Comput. VBC'96*, pp. 439–448. Hamburg.
- Zilles, K., Schlaug, G., Matelli, M., Luppino, G., Schleicher, A., Qu, M., Dabringhaus, A., Seitz, R. & Roland, P. 1995 Mapping of human and macaque sensorimotor areas by integrating architectonic, transmitter receptor, MRI and PET data. *J. Anat.* **187**, 515–537.
- Zilles, K. (and 11 others) 1997 Quantitative analysis of sulci in the human cerebral cortex: development, regional heterogeneity, gender difference, asymmetry, intersubject variability and cortical architecture. *Hum. Brain Mapp.* **5**, 218–221.