

Measuring Cortical Thickness

Jason Lerch



Department of Neurology and Neurosurgery
McGill University
Montreal, Canada

July 2001

A Masters of Science Thesis Proposal.

Chapter 1

Introduction

The overall goal of this master's project is to study the effect of aging on the thickness of the cerebral cortex. There are three inherent challenges in this project:

- Creating a coherent definition of cortical thickness.
- Implementing that definition numerically in the context of real MRI data.
- Applying the thickness implementation to a real dataset of an aging population, and interpreting the results.

Defining cortical thickness is a surprisingly non-trivial task. The question is moreover a fundamental one: what is it, ideally, that one wants to measure? I will try to argue that previous approaches of drawing a straight line from the surface to the white matter are inadequate, and that thickness is not a property of just the surface, but instead should be definable at any point within the cortical mantle, and that a thickness metric ought to try to model the functional unit of the cortex, the cortical column.

The numerical/computational implementation of any definition of cortical thickness raises further issues. There are the inherent limitation of MRI to consider, such as limited resolution and partial volume effect. Important also are implementation issues involving concerns such as choice of algorithms and use of system resources.

Lastly, the more theoretical and computational issues also have to be applied to a practical goal: the study of how cortical thickness changes with age. We know from previous MRI studies that there is a noticeable decline in the cortical mantle with age, though the cellular mechanism that underlie that decline still need to be elucidated. Animal studies appear to suggest that it is not a

decline in number of neurons, but rather a reduction in arborisation and shrinkage of the neurons themselves. The resulting effect, however, is that the thickness of the cortex declines, the sulci widen and the gyri narrow. A study of changes in cortical thickness thus provides an interesting neurological challenge along with a simple statistical model against which to test measures of cortical thickness.

Chapter 2

Measuring Cortical Thickness

The cerebral cortex has been called the “highest achievement of biological evolution and the neural substrate of human mental abilities” [26]. The cortex has a surface area of on average 2.5 square feet, with a normal thickness of about 3mm [4, 11, 17]. It is a highly convoluted structure, the degree of folding likely related to an evolutionary need to increase surface area without a corresponding increase in intracranial size [17]. Overall, the cortex is estimated to contain 14 billion neurons, principally of the pyramidal, stellate, and fusiform varieties [4].

Measuring cortical thickness is an important task for both normal and abnormal neuroanatomy. The cortical mantle varies in thickness depending on the region of the cortex, with considerable variation between individual brains as well as between hemispheres of the same brain [12]. In normal brains the cortex tends to be thinnest in the calcarine cortex at around 2mm and highest in the precentral gyrus at around 4mm [11]. Thickness information is thus both interesting in its own right as well as a useful aid in such tasks as sulcal labelling [36, 19, 10, 27].

In pathological cases cortical morphology has been known to vary in epilepsy [11], mental retardation [12], Schizophrenia [11], anorexia nervosa [11], and Alzheimer’s disease [11, 30]. Moreover, it has been shown that there is a general pattern of cortical evolution with age consisting of widening of the sulci along with a thinning of the cortical mantle [17]. These studies clearly suggest that morphological changes in the thickness of the cortex are associated with meaningful functional differences across groups.

The argument that I would like to lay out below is that the cerebral cortex should not be considered as a homogenous and indistinguishable mass of grey matter, but instead ought to be treated in terms of its laminar and columnar functional organisation. Ultimately the question of changes

in cortical thickness must be related back to changes in the functional output of the cortex, and as such the thickness metric ought to reflect the geometric organisation of these functional units as closely as possible.

2.1 Cortical Thickness: An Anatomical Definition

The cerebral cortex is organised along the following three principles:

1. It is organised in vertical layers, each layer containing different cellular organisation and synaptic connections.
2. It is organised in vertical columns coding for a single function and stretching across all layers.
3. These columns are further often organised into *hypercolumns* which combine all of the different functional columns for one area (such as one part of the visual field).

The cortex has a laminar organisation, organised into six separate layers throughout the neocortex, fewer in the allocortex [12, 18]. The layers are (from the surface towards the white matter): (I) the molecular layer, (II) the corpuscular layer, (III) the pyramidal layer, (IV) the granular layer, (V) the ganglionic layer, and (VI) the multiform layer [3]. The differences between these layers is mostly distinguished on the basis of pyramidal cell impregnating staining techniques which reveal different packing densities of pyramidal cells in the various laminae [3].

The anatomical layers of the neocortex furthermore also have distinct characteristic synaptic connections. Layer $4C\alpha$ of the striate cortex, for example, receives input from the magnocellular layers of the Lateral Geniculate Nucleus of the Thalamus, sends projection interneurons to layer 4B, from there to layers 2 and 3 which in turn project to extrastriate areas [33].

The functional unit of cortical processing is organised in a columnar fashion. The smallest processing unit of the mature cortex is the *minicolumn*, which extends approximately vertically across layers 2-6 and contain all major phenotypes of cortical neurons [20]. The minicolumns are further bound together through dense short-range horizontal connections into *cortical columns* or *modules* [20]. The cortical columns are unified by certain static or dynamic properties, a classic example being the orientation columns of the striate cortex [20, 33].

Evidence for the functional organisation comes from microelectrode penetrations into the cortex along with transection and nerve regeneration studies. Electrode penetration studies of the

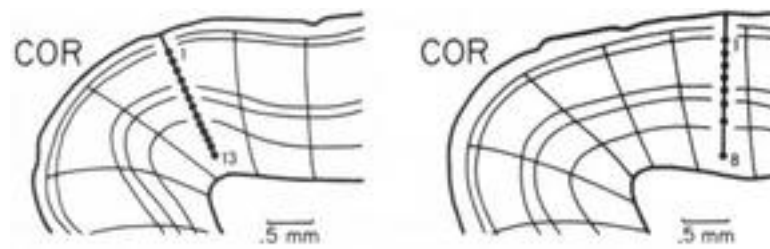


Fig. 2.1 Thin lines vertical lines illustrate the cortical columns, horizontal lines represent layers. Taken from [20]

striate cortex found that perpendicular penetrations showed constant responses to a stimulus of a single orientation, whereas penetrations made nearly parallel to the surface found a consistent change in sensitivity to differently oriented stimuli [20, 22]. Figure 2.2 provides a schematic illustration of these results.

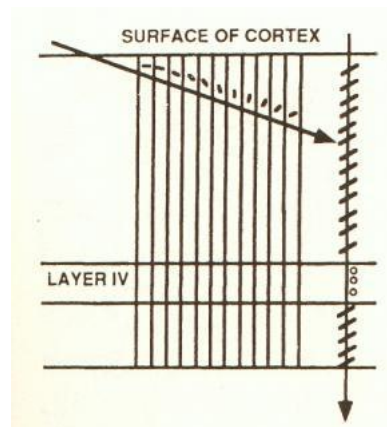


Fig. 2.2 This figure shows orientation sensitivity of columns in the striate cortex. Vertical penetrations show a single orientation, whereas parallel penetrations show multiple orientations. From [20]

Similarly, studies in the somatic sensory cortex found columns coding for both modality (i.e. $A\alpha$ fibers from Meissner receptors) as well as location (point on finger). Moreover, nerve-regeneration studies found that in the non-transected animal these columns would form a smooth transition, whereas after transection and resuture of the contralateral medial nerve the columns are separated by 50-60 μm [20].

The geometric organisation of these columns are critical for the purposes of measuring cortical thickness. Schematically, columns run perpendicular to the pial surface - but closer examination reveals that this is not truly the case. Instead, the columns appear to find the shortest path from

one surface to the other while taking the curvature of the two surfaces into account. Columns thus do curve in their passage from the white matter through the cortical layers. This geometric organisation of the cortex can be seen in two slice preparations from figure 2.3. Moreover, figure 2.1 further illustrates this geometric property schematically. Lastly, it is important to note that each column codes for one modality or series of complementary modalities (such as ocular dominance and orientation sensitivity) at a certain location, and that therefore theoretically direct overlap of these columns is not possible.

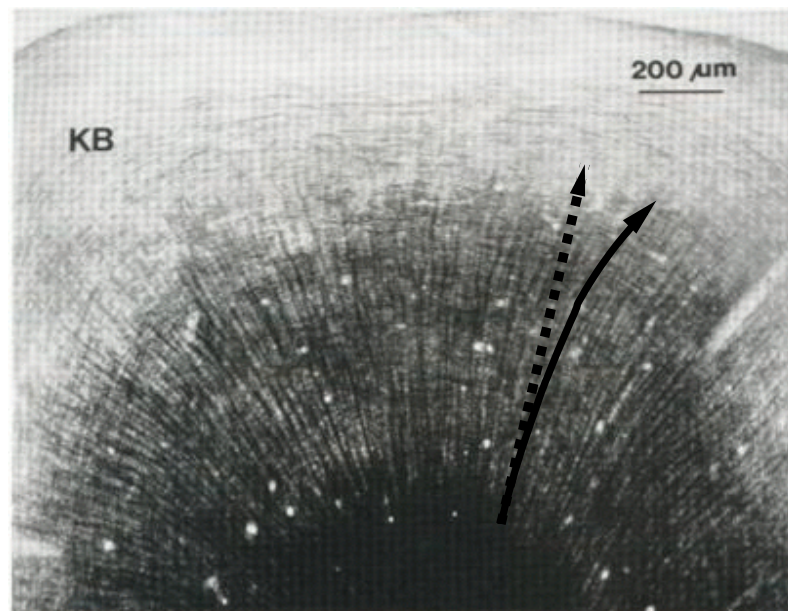


Fig. 2.3 A myelin stained section of the cortex. It is evident here that the myelin fibers curve in their progression towards the outer layers. The dotted black line shows what a straight measurement of thickness would look like compared to the filled line which traces the actual orientation of the myelin fibers. Taken from [3].

The inherent organisation of the cerebral cortex defined above leads to the hypothesis that the most significant functional change in the physiology of the cortex would consist of changes along the trajectory of the cortical columns. The organisation of the cortex further implies that any point within the cortical mantle belongs to a single cortical column. If cortical thickness is defined as changes along the axes of the cortical columns then one logical implication is that thickness is the property of any place within the cortical mantle. Moreover, a precursory examination of (especially myelin) stained sections of the cortex reveals that the columns do not necessarily follow a direct line from the white to the CSF surface, instead the directional angle appears to at least have

a rough relation to the layering of the cortical laminae.

2.2 A Cortical Thickness Metric: Literature Review and an operational definition

The review of cortical anatomy above naturally leads to a series of requirements for a definition of cortical thickness - these requirements will, however, have to be tempered with the limitations inherent in current Magnetic Resonance acquisition protocols. On the basis of cortical anatomy, I would propose that any implementation of cortical thickness should have the following properties:

1. Thickness is the property of the entire cortical mantle. Thus, any point within the cortical mantle should have one and only one thickness associated with it.
2. Thickness is the property of the functional unit of the cortex, the cortical column.
3. The thickness measurement at any one point ought to be the shortest distance that meets the above criteria.

In order for the points above to be implemented the nature of the dataset to which any definition will be applied must be considered first. In this case that data will come from Magnetic Resonance Imaging. The first and most obvious limitation is the inherent resolution of an MRI volume: the best one can expect in a real dataset is one millimetre isomorphic sampling, though it might be possible to acquire 0.5 millimetre datasets for verification purposes. Even at 0.5 millimetres, however, individual neurons are clearly not visible and even the cortical layers are next to undifferentiable. The cortex in an MR image thus becomes a seemingly homogeneous grey mass, the finer points of anatomy that were delineated above disappear. The task of a thickness metric is thus to mathematically approximate those features in the absence of them actually being visible.

This then leads to a functional definition of thickness as measurable in MRI:

1. Thickness is the property of every voxel in the cortical mantle, and every voxel in the cortex should have one and only one thickness associated with it. This assertion has several implications:
 - (a) Thickness will be a volumetric measurement.

- (b) Lines of thickness between the two surfaces may not intersect.
- 2. The measurement of cortical thickness should take a priori knowledge about cortical morphology into account and try to approximate that information where possible.
- 3. The thickness measurement at any one point ought to be the shortest distance that meets the above criteria.

2.2.1 Cortical Thickness Metrics in Literature

There have been several attempts made at measuring cortical thickness, including both post-mortem studies as well as computational studies using MRI. With only a single exception, they have used a variation of the what I shall term *straight-line approach* to measuring cortical thickness. In short, this approach finds the shortest line from the cortical surface to the grey and white matter boundary - though the direction which that line could take may be constrained. The use of this straight-line method was initially implemented in the study of post-mortem specimen, where the investigator would either insert a probe through the outer surface and measure the distance along the (rather haphazardly chosen) angle of the probe towards the white matter, or else the investigator would examine a slice of cortex and use a jeweller's eyepiece to measure the distance between the white matter and the surface along the angle of the slice cut [16].

There is an obvious and clear limitation in the post-mortem studies described above, namely the choice of angle along which to measure. The highly convoluted nature of the cortex makes this choice a tricky task indeed. Ultimately, the outcome will over-estimate the thickness at any one location unless the slice or probe penetration angle is perfectly orthogonal to the cortical surface. There is, moreover, another inherent problem in post-mortem data analysis, namely the issue of shrinkage of the specimen, which can clearly lead to incorrect absolute values for the thickness at any one point [12].

Studies which measure cortical thickness from MR images have been few and far between, most likely due to the inherent difficulty in executing the task correctly. There have been two types of approaches used: one which attempts to replicate the jeweler's eyepiece type measurement by measuring the distance from the surface to the white matter on a slice. The other approach tries to separate the two surfaces (grey/cortico-spinal fluid (CSF) and grey/white) and create object representation of these two surfaces only to then find the closest point on one surface given a point on the other.

The first approach faces the same problem as the post-mortem studies: picking the correct slice angle along which to measure the thickness at any one point. That is a very difficult task, made even more difficult by the fact that MRI is discrete data rarely sampled higher than one millimetre. Moreover, it is also a very labour intensive operation, making this technique prohibitive for use in larger studies. One example of the application of this method comes from the work of Meyer et al. who measured the thickness of the cortex perpendicular to the central sulcus using a jeweler's eyepiece - the goal of this exercise being the validation of thickness measurements as a technique of identifying the central sulcus [19].

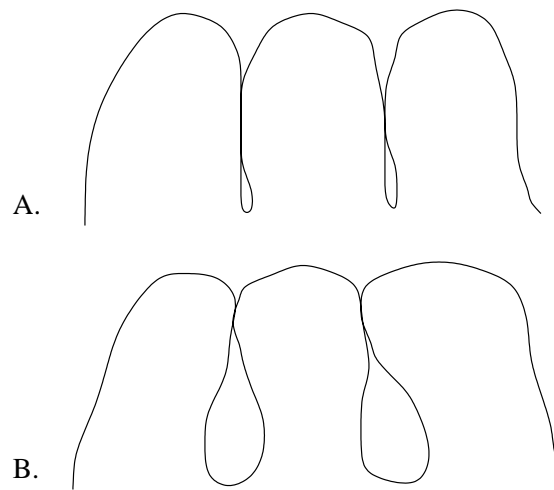
The second approach uses a much more sophisticated series of data processing techniques, and does to considerable extent eliminate the problems outlined for the jeweler's eyepiece. This type of approach has to deal with three main issues:

1. Segmentation of the MR volume into its component tissue types (usually white matter, grey matter, cortico-spinal fluid, and background).
2. Separation of the gyri that have been fused through the partial volume effect.
3. Construction of the actual surfaces into some polygon model.

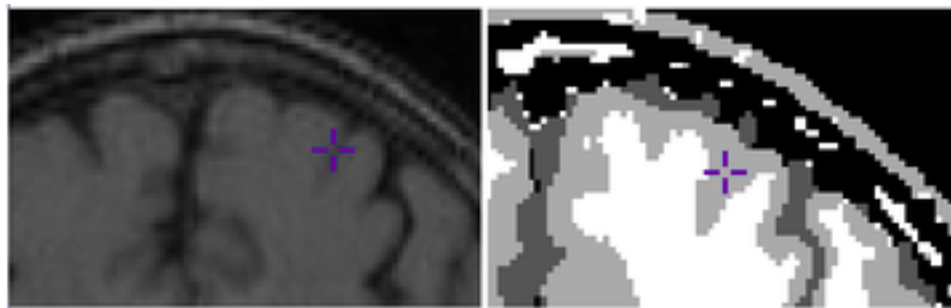
Tissue segmentation is a fairly common technique in MR image processing that has been discussed in depth elsewhere (c.f. [29, 34, 35]) and will therefore be left behind in the discussion. Suffice it to say that the various techniques generally are capable of a solid, though not perfect, separation of cortical grey matter from the underlying white matter.

The partial volume effect - also often referred to as the problem of the buried cortex - induced fusing of gyri, on the other hand, is a much more difficult issue. It is also likely the reason why so few MRI studies of thickness or other cortical measurements have been attempted. The problem is the following: due to the inherent coarseness of sampling in MRI adjacent gyri can overlap in the same voxel. This causes the output image to have an intensity value at that location that causes the classifying algorithm to treat that voxel as grey matter. The end result is that two gyri which in reality are separated by cortical CSF now appear fused in the output image. Any algorithm which simply tries to extract the cortex by marching along the CFS grey matter boundary will therefore miss the cortical folding that has been obscured by the partial volume effect artefact - see figure 2.4 for an illustration.

A survey of the literature reveals three distinct types of methods for dealing with the problem of the buried cortex. The first of these can be seen in the implementation of Magnotta et al, who



(a) Illustration



(b) MRI example

Fig. 2.4 An illustrative example of the buried cortex phenomenon. Part A shows gyri abutting each other, which would result in a fused grey matter area in MRI. Part B shows gyral crowns abutting each other, which would fuse only those two parts. Partly adapted from [17]. The last image shows an actual example from an MR image, the T1 volume on the left and the classified output on the right.

circumvent the problem through eroding of the cortical grey matter in a consistent fashion which opens up the sulci [17]. A second approach is advocated by Jones et al, who use a combination of edge thinning and gradient information to find the obscured sulci [11]. The last method, developed both by MacDonald et al and Fischl and Dale, uses anatomical constraints derived from the white matter surface to deform the grey matter surface onto the correct target [16, 8, 15, 6]. One of the differences between the various methods is that the first two clearly separate the removal of the partial volume effect from the actual creation of the surface (which happens afterwards) [17, 11], whereas the last method tries to circumvent the buried cortex problem as part of the creation of the surfaces [16, 8, 15, 6].

Once the surfaces exist the thickness at any one location has to be found. The most common approach - used by everyone except Jones et al - is some variation on the straight line measurement. The technique used involves picking a point on one of the surfaces (grey or white matter surface) and finding the closest point on the other surface [8, 17, 16]. The variations on this theme are that the choice of point can be constrained to follow the surface normal [16], or that the starting point is not one of the surfaces but instead the midline between the two surfaces (which is more of an artefact of the cortical thinning algorithm used to counter the partial volume problem) [17].

The straight line approach is inherently problematic, however. In its essence, the issues are threefold:

- Thickness is defined only at the surface and not throughout the mantle.
- Thickness measurements will vary depending on which surface you measure from.
- Lines of thickness can intersect.

The first point deals with where one defines cortical thickness. The straight line approach measures thickness only at one of the surfaces of the cortex, and the results will furthermore vary depending on which surface you start out from. See figure 2.5 for an illustration of how this type of measurement can result in multiple thickness measurements for one point in the surface.

The problem is more fundamental than that, however. The question one needs to ask when measuring cortical thickness is what anatomic substrate is it that is under consideration. This becomes especially relevant when the question under investigation is related to relative change over time or between groups: what are the cellular processes that are being modeled by measuring changes in cortical thickness? I hope to have convincingly argued that these processes are

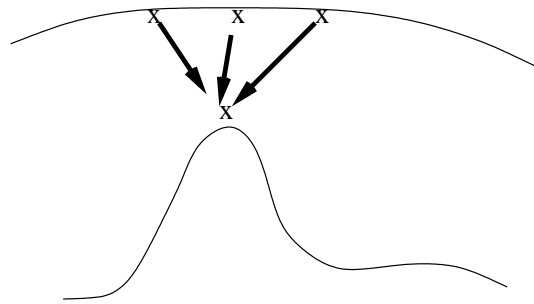


Fig. 2.5 This illustration shows how one point on the inside surface can have multiple thickness measurements associated with it if a straight-line measurement from the outer surface is used.

fundamentally related to the functional unit of the cortex, the cortical column. While it is not possible for MRI to directly investigate these columns, we can take several properties of them into consideration. One of the most fundamental of these properties is the mutual exclusivity of the cortical columns, and attribute that is inherent in the electrical and chemical signalling properties of neurons. So if, when modelling changes of thickness we aim to at least approximate columnar change, then any one point in the cortical mantle can have one and only one thickness associated with it. It is clear that any one voxel in an MRI image will contain multiple columns and that the thickness measurement thus will never estimate the length of such a column, but instead approximate their geometric orientation. The straight-line based approach fails on that account, as it measures thickness only from the surfaces and not at any point in the mantle, and moreover the lines of thickness between the surfaces can intersect.

The thickness metric outline by Jones et al circumvents all of these problems by borrowing an approach from computational fluid dynamics [11]. In outline form, the cortical thickness is computed through a partial differential equation, *Laplace's equation*. This is a second-order partial differential equation for a scalar field Ψ that is enclosed between two surfaces and takes the mathematical form:

$$\Delta^2\Psi = \frac{\partial^2\Psi}{\partial x^2} + \frac{\partial^2\Psi}{\partial y^2} + \frac{\partial^2\Psi}{\partial z^2} = 0$$

This is a harmonic function or potential function, and has the property - ideal for the measurement of cortical thickness - of effectively describing a nested series of surfaces that make a smooth transition between the inner and outer surface [11]. Field lines between the two surfaces can then be computed, and the thickness at any one point in the mantle can be found by integrating these

streamlines [11]. See figure 2.6.

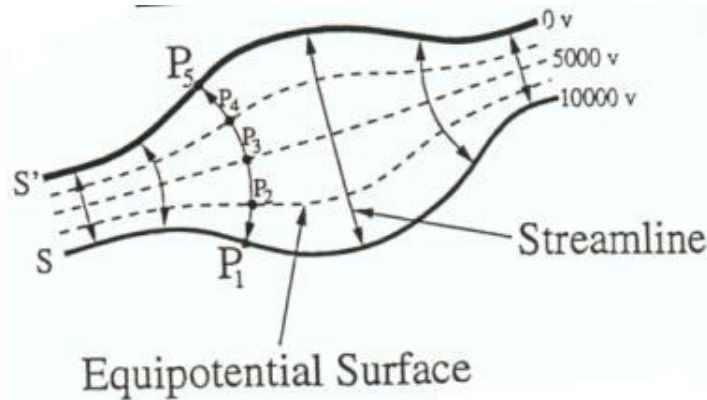


Fig. 2.6 A two-dimensional example of Laplace's method. Each surface is assigned a set value and intermediate surfaces are created through solving of the partial differential equation. Field lines can then be created which represent the thickness at that point. From [11].

The Laplacian approach advocated by Jones et al circumvents the problems of the straight-line approach. Does it directly model the structure of the cortical columns? The answer is clearly no. We do not have the information in MRI to be able to carry that task out. What the laplacian method does do, however, is conserve some of the properties that columns are presumed to have (the most important being mutual exclusivity and definition throughout the cortical mantle), all the while maintaining a geometric approach that is an intuitively sensible approximation of the columnar layout.

2.2.2 Suggested Cortical Thickness Implementation

It is now time to outline the cortical thickness metric that I would implement. In essence it will combine the strengths of MacDonald et als cortical extraction [15, 16] with the fluid dynamics approach outline by Jones and colleagues [11]. Figure 2.7 illustrates the suggested processing steps.

The initial steps in the processing pipeline are well known and have been used in multiple studies originating from the Brain Imaging Centre (c.f. [23, 32]). It proceeds by removing field inhomogeneities from the native data [28] while simultaneously finding the transformation matrix from native space to MNI Talairach space [5]. The nonuniformity corrected volumes are then resampled using the Talairach transformation [5]. This is followed by tissue classification [34, 35]

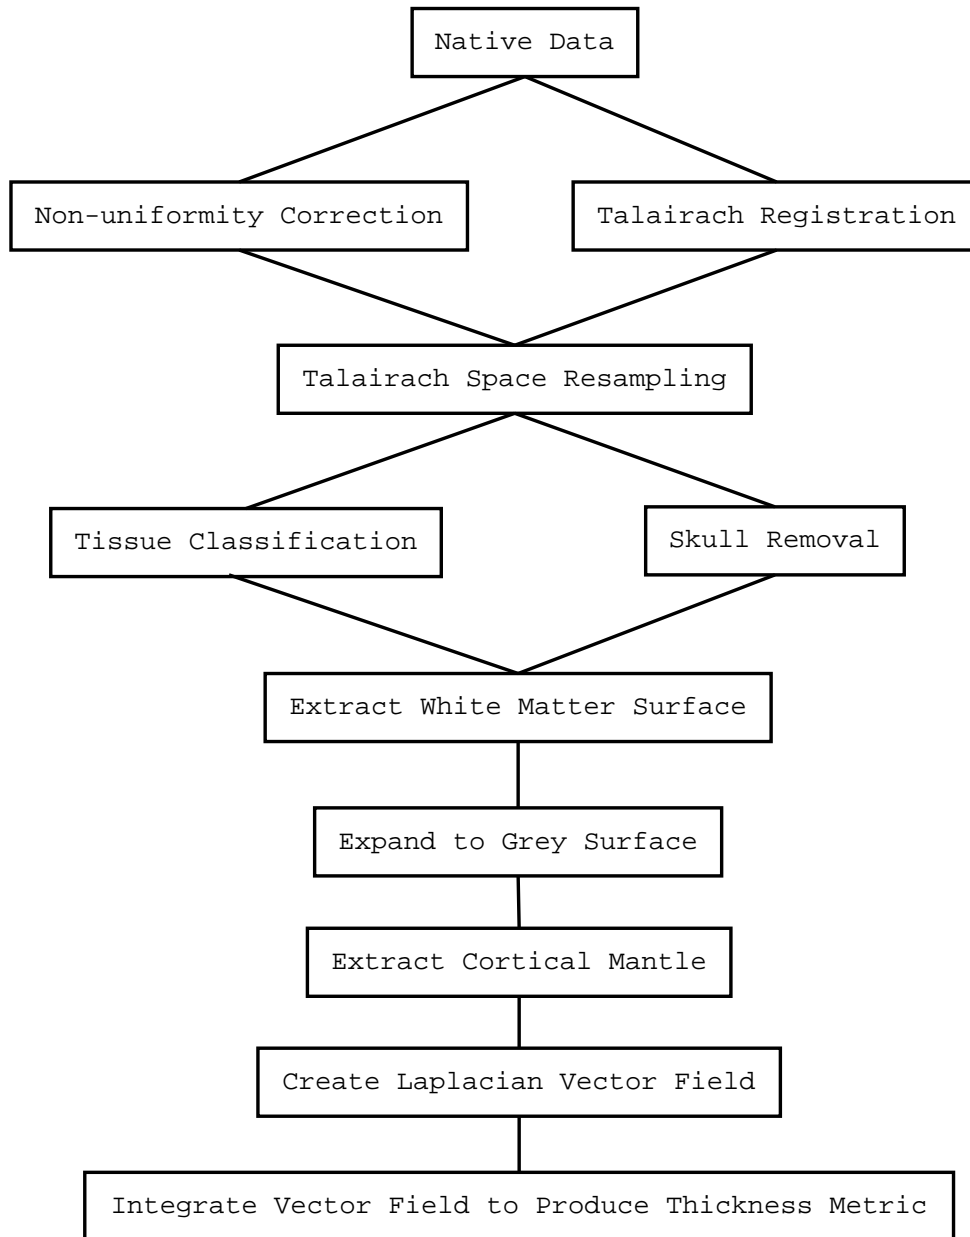


Fig. 2.7 A flowchart illustrating the processing steps for implementing a Laplacian thickness metric.

and initial removal of the skull [15].

The following two steps are rather compute intensive - taking about 35 hours on an SGI Origin 200 - and proceed by first deforming a sphere to the white matter surface (as identified in the classification step), and then expanding that surface to the grey/CSF boundary [15, 16]. The fact that the surfaces retain the inherent topology of a sphere is inherently advantageous for the measurement of cortical thickness, as it allows for easy inflating of the surface later on to visualise the buried sulci.

Once the surfaces exist the volumetric cortical mantle has to be extracted. This can be done through a computational geometry approach, where a few points that lie between the surfaces are found and the rest of the mantle is then iteratively filled in. The cortical mantle is then used to initialise the solver of the boundary-value problem partial differential equation. The mantle itself is set to a neutral value, the outer surface to 10,000 and the inner surface to 0. The equation is then iteratively solved through Jacobian relaxation [11].

One issue that does arise here is the sampling of the volume over which Laplace's equation is to be solved. If left at one millimetre sampling, the grid that encompasses the cortex will never be more than 5 voxels thick, implying that the vector field will not contain sufficient information for an intelligent solution. It is thus preferable to use a finer grid, which of course raises issues of memory consumption. The initial approach taken will be to subsample the dataset to 0.5 millimetres, which can hopefully be refined to an even smaller sampling later on through the use of sparse matrices to reduce usage of system resources.

The final step will then be to normalise the vector field and integrate it at each voxel in order to determine the thickness at that voxel [11]. The thickness information will then be available volumetrically, but can also be fused back on to the surfaces for surface based visualisation.

Chapter 3

The Aging Cortex

The most interesting results likely to emerge from applications of a cortical thickness metric are relative changes in longitudinal studies or changes between groups in cross-sectional studies. This section will therefore briefly examine some changes that can be expected in an aging population.

Studies of aging and the brain have revealed a dichotomy of findings - MRI based studies have found a large decline associated with age in the grey matter (c.f. [24, 31]), whereas more molecular studies have shown that little or no neuronal loss with age (c.f. [2]).

The answer appears to be that the cellular changes that occur with age are related to a reduction in integrity of the myelin fibers in the cortex [21] along with a decline in dendritic arborisation and spine density [2]. A part of the differences in the two types of aging studies can also be attributed to species differences, as most MRI studies are performed on human subjects, whereas the slice preparations usually involve animal models.

It is clear, however, that MRI reveals a significant age related loss of grey matter tissue, and that moreover that loss is spread throughout the entire cortex. Recent data where 806 subjects from the Japanese Aging Project were examined using voxel-based morphometric techniques show this global decline aptly - see figure 3.1 for an illustration. Other studies focusing on specific brain regions have also shown considerable age-related decline (c.f. [31, 24]). Part of the question that arises is what the decline seen in MRI represents. While this question is still far from being answered, general shrinkage of neurons through reduction in arborisation along with a loss of supporting structures like myelin sheaths probably accounts for a larger part of this decline than neuronal death.

A few other studies have examined cortical changes with normal and pathological aging as

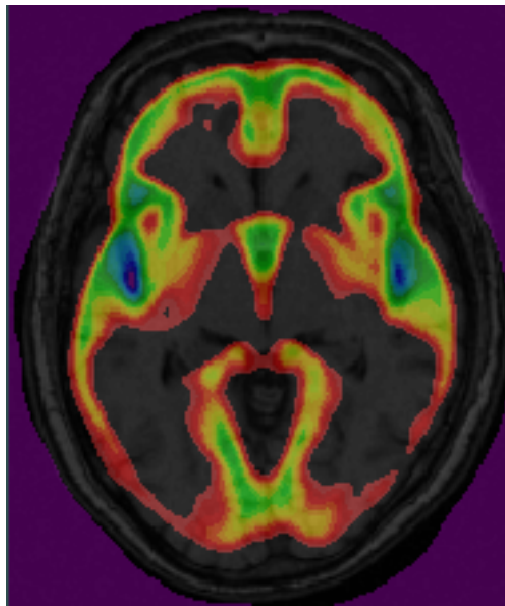


Fig. 3.1 Voxel-based T-statistics from the Japanese Aging Project. T values range from -10 to -32.

well. Magnotta et al examined changes in sulcal and gyral pattern along with cortical thickness (using a variant of the straight-line method, unfortunately). What they found was that the gyri become more steeply curved with age, the sulci widen and cortical thickness declines [17]. These changes are more dominant in males than females - a fact that seems to hold for general age related decline in the brain (c.f. [24]).

Analysing an aging set using a cortical thickness metric provides several advantages over the traditional approaches. MRI studies have, for the most part, employed two separate techniques. The first is the manual segmentation of structures of interest by trained neuroanatomists - a technique which can provide highly accurate results but is very labour intensive and has to contend with problems of inter and intra rater reliability. The second technique uses voxel based morphometry (VBM) [23, 32, 1] where the density of a tissue type at each location is examined. VBM is fully automated and thus a powerful way to examine tissue changes across a statistical model. When one is examining cortical changes, however, measuring the thickness of the cortex provides a closer approximation to the underlying anatomical reality than voxel density as measured by VBM. In other words, the tissue matter maps used by VBM are only useful within a statistical context, whereas an individual thickness map can already tell the investigator something about the subject (i.e. focal thickness increases in epilepsy).

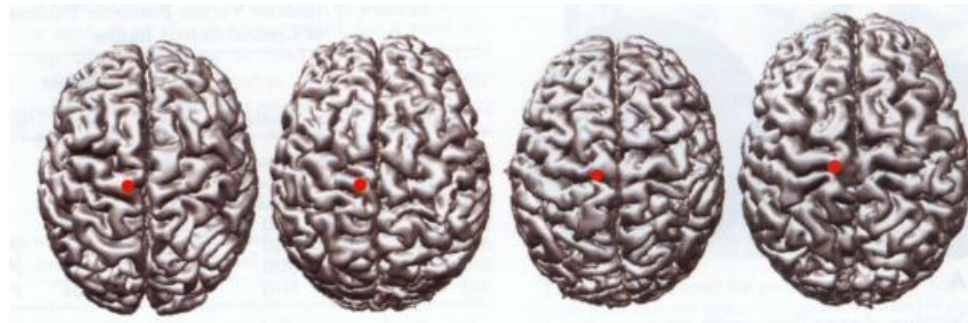


Fig. 3.2 Rough correspondance of vertices shown by marking the same vertex in four different subjects. From [16]

One of the as of yet unresolved issues that will arise in the study of cortical thickness and aging is that of statistical analysis. Ideally one would be able to generate parametric maps of thickness changes across populations in a similar vein as voxel-based morphometric techniques do. The difficulty, however, lies in the issues of accurate registration of cortical features and the blurring of resulting data to account for inter-subject variability. Registration of cortical features is difficult due to the large inherent variability in cortical folding patterns between individuals (c.f. [9, 13, 14, 7, 25, 36]). The generation of cortices to be used will be at least roughly registered, as the procedure takes place in Talairach space. Moreover, since the surfaces are deformed from the same sphere, each vertex will be in a roughly corresponding location (see figure 3.2) for each of the subjects [15, 16], and will therefore allow at least some level of comparison until better registration methods can be developed.

Chapter 4

Conclusions

I hope to have outlined a coherent argument for the implementation of the Laplacian cortical thickness measurement. There will, of course, be many challenges ahead in implementing the above proposal, though it should be doable within the allotted time of a masters project.

The current implementation of the processing pipeline is nearly complete. To refer back to figure 2.7, everything except for the integration of the vector field has been completed, though surely there are issues still to be resolved in those steps as well. Integration of a vector field is a fairly common task, and should also not take too long to complete.

The dataset that will initially be used will be the Sendai dataset of 1000 plus subjects from the Japanese Aging Project. Concurrently some validation studies will also have to be done. This will not be easy, however, since the approach to measuring thickness that will be implemented is different from those performed in previous post-mortem studies, and thus potentially not directly comparable (though the results ought to still roughly correspond). One potential verification will involve testing the output of the fully automated pipeline against thickness analysis of manually segmented gyri. Other verification could potentially use high-resolution MRI or cryosection data where cortical layering is visible.

There are also many other datasets and projects available where the application of thickness measurements could prove interesting. Among them are datasets from epilepsy, pediatric development, and Alzheimer's and Mild Cognitive Impairment, just to mention a few.

Cortical thickness measurements in vivo have only recently become possible, and have never been carried out on large datasets. This project can thus provide both a useful definition of cortical thickness along with an effective implementation. Moreover, this technique applied to datasets

such as aging can prove quite valuable for the field of neuroscience.

References

- [1] John Ashburner and Karl Friston. Voxel based morphometry - the methods. *NeuroImage*, 11:805–821, 2000.
- [2] J.C. Baron and G. Godeau. Human aging. In Arthur W. Toga and John C. Mazziotta, editors, *Brain Mapping: The Systems*. Academic Press, 2000.
- [3] Heiko Braak. Architectonics as seen by lipofuscin stains. In Alan Peters and Edward G. Jones, editors, *Cerebral Cortex: Cellular Components of the Cerebral Cortex*, volume 1. Plenum Press, 1984.
- [4] Malcolm B. Carpenter. *Core Text of Neuroanatomy*. Williams & Wilkins, third edition, 1985.
- [5] D. L. Collins, P. Neelin, T. M. Peters, and A. C. Evans. Automatic 3d intersubject registration of mr volumetric data in standardized talairach space. *Journal of Computer Assisted Tomography*, 18:192–205, 1994.
- [6] Anders M. Dale, Bruce Fischl, and Martin I. Sereno. Cortical surface-based analysis i: Segmentation and surface reconstruction. *NeuroImage*, 9:179–194, 1999.
- [7] Heather A. Drury, David C. Van Essen, Maurizio Corbetta, and Abraham Z. Snyder. Chapter 19: Surface-based analyses of the human cerebral cortex. In Arthur W. Toga, editor, *Brain Warping*. Academic Press, 1999.
- [8] Bruce Fischl and Anders M. Dale. Measuring the thickness of the human cerebral cortex from magnetic resonance images. *PNAS*, 97(20):11050–11055, 2000.
- [9] Bruce Fischl, Martin I. Sereno, and Anders M. Dale. Cortical surface-based analysis ii: Inflation, flattening, and a surface-based coordinate system. *NeuroImage*, 9:195–207, 1999.
- [10] Georges Le Goualher, Anne Marie Argenti, Michel Duyme, William F.C. Baare, H.E. Hulshoff Pol, Christian Barillot, and Alan C. Evans. Statistical sulcal shape comparisons: Application to the detection of genetic encoding of the central sulcus shape. *No Idea what journal this is from*, 2000.

-
- [11] Stephen E. Jones, Bradley R. Buchbinder, and Itzhak Aharon. Three-dimensional mapping of cortical thickness using laplace's equation. *Human Brain Mapping*, 11:12–32, 2000.
- [12] Noor Kabani, Georges Le Goualher, David MacDonald, and Alan C. Evans. Measurement of cortical thickness using an automated 3-d algorithm: A validation study. *NeuroImage*, 13:375–380, 2001.
- [13] D.N. Kennedy, N. Lange, N. Makris, J. Bates, J. Meyer, and V.S. Caviness Jr. Gyri of the human neocortex: An mri-based analysis of volume and variance. *Cerebral Cortex*, 8:372–384, jun 1998.
- [14] Gabriele Lohman, D. Yves von Cramon, and Helmuth Steinmetz. Sulcal variability of twins. *Cerebral Cortex*, 9:754–763, oct-nov 1999.
- [15] David MacDonald. *A method for identifying geometrically simple surfaces from three dimensional images*. PhD thesis, School of Computer Science, McGill University, 1997.
- [16] David MacDonald, Noor Kabani, David Avis, and Alan C. Evans. Automated 3-d extraction of inner and outer surfaces of cerebral cortex from mri. *NeuroImage*, 12:340–356, 2000.
- [17] Vincent A. Magnotta, Nancy C. Andreasen, Susan K. Schultz, Greg Harris, Ted Cizadlo, Dan Heckel, Peg Nopoulos, and Micheal Flaum. Quantitative in vivo measurement of gyri-fication in the human brain: Changes associated with aging. *Cerebral Cortex*, 9:151–160, mar 1999.
- [18] John H. Martin. *Neuroanatomy Text and Atlas*. Appleton & Lange, 2 edition, 1996.
- [19] Joel R. Meyer, Sudipta Roychowdhury, Eric J. Russel, Cathy Callahan, Darren Gitelman, and M. Marsel Mesulam. Location of the central sulcus via cortical thickness of the precentral and postcentral gyri on mr. *American Journal of Neuroradiology*, 17:1699–1706, oct 1996.
- [20] Vernon B. Mountcastle. *Perceptual Neuroscience: The Cerebral Cortex*. Harvard University Press, 1998.
- [21] Kirsten Nielsen and Alan Peters. The effects of aging on the frequency of nerve fibers in rhesus monkey striate cortex. *Neurobiology of Aging*, 21:621–628, 2000.
- [22] K. Obermayer and G.G. Blasdel. Geometry of orientation and ocular dominance columns in monkey striate cortex. *Journal of Neuroscience*, 13(10):4115–4129, 1993.
- [23] Tomas Paus, Alex Zijdenbos, Keith Worsley, D. Louis Collins, Jonathan Blumenthal, Jay N. Giedd, Judith L. Rapoport, and Alan C. Evans. Structural maturation of neural pathways in children and adolescents: In vivo study. *Science*, 283:1908–1911, march 1999.

- [24] J.C. Pruessner, D.L. Collins, M. Pruessner, and A.C. Evans. Age and gender predict volume decline in the anterior and posterior hippocampus in early adulthood. *Journal of Neuroscience*, 27(1):194–200, 2001.
- [25] J. Rademacher, V.S. Caviness Jr., H. Steinmetz, and A.M. Galaburda. Topographical variation of the human primary cortices: Implications for neuroimaging, brain mapping, and neurobiology. *Cerebral Cortex*, 3:313–329, jul-aug 1993.
- [26] John L.R. Rubenstein and Pasko Rakic. Genetic control of cortical development. *Cerebral Cortex*, 9:521–523, sep 1999.
- [27] Francisca Aina Sastre-Janer, Jean Regis, Pascal Belin, Jean-Francois Mangin, Didier Dormont, Marie-Cecile Masure, Philippe Remy, Vincent Frouin, and Yves Samson. Three-dimensional reconstruction of the human central sulcus reveals a morphological correlate of the hand area. *Cerebral Cortex*, 8:641–647, oct-nov 1998.
- [28] J. G. Sled, A.P. Zijdenbos, and A.C. Evans. A nonparametric method for automatic correction of intensity nonuniformity in mri data. *IEEE Transactions on Medical Imaging*, 17:87–97, 1998.
- [29] Rik Stokking, Koen L. Vincken, and Max A. Viergever. Automatic morphology-based brain segmentation (mbrase) from mri-t1 data. *NeuroImage*, 12:726–738, 2000.
- [30] P.M. Thompson, J. Moussai, S. Zohoori, A. Goldkorn, A.A. Khan, M.S. Mega, G.W. Small, J.L. Cummings, and A.W. Toga. Cortical variability and asymmetry in normal aging and alzheimer’s disease. *Cerebral Cortex*, 8:492–509, sep 1998.
- [31] D.J. Tisserand, P.J. Visser, M.P.J. van Boxtel, and J. Jolles. The relation between global and limbic brain volumes on mri and cognitive performance in healthy individuals across the age range. *Neurobiology of Aging*, 21:569–576, 2000.
- [32] K.E. Watkins, T. Paus, J.P. Lerch, A. Zijdenbos, D.L. Collins, P. Neelin, J. Taylor, K.J. Worsley, and A.C. Evans. Structural asymmetries in the human brain: a voxel-based statistical analysis of 142 mri scans. *Cerebral Cortex*, In Press, 2001.
- [33] Robert H. Wurtz and Eric R. Kandel. Central visual pathways. In Eric R. Kandel, James H. Schwartz, and Thomas M. Jessel, editors, *Principles of Neural Science*. McGraw Hill, fourth edition, 2000.
- [34] Alex Zijdenbos, Alan Evans, Farhad Riahi, John Sled, Hing-Cheung Chui, and Vasken Kollokian. Automatic quantification of multiple sclerosis lesion volume using stereotaxic space. In Karl-Heinz Höhne and Ron Kikinis, editors, *Proceedings of the Fourth International Conference on Visualization in Biomedical Computing (VBC)*, pages 439–448, Hamburg, Germany, 1996. Springer.

-
- [35] Alex Zijdenbos, Reza Forghani, and Alan Evans. Automatic quantification of ms lesions in 3D MRI brain data sets: Validation of insect. In *Proceedings of the First International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 439–448, Cambridge MA, USA, October 1998.
- [36] Karl Zilles, Axel Schleicher, Christian Langeman, Katrin Amunts, Patricia Morosan, Nicola Palomero-Gallagher, Thorsten Schormann, Hartmut Mohlberg, Uli Burgel, Helmut Steinmetz, Gottfried Schlaug, and Per E. Roland. Quantitative analysis of sulci in the human cerebral cortex: Development, regional heterogeneity, gender difference, asymmetry, inter-subject variability and cortical architecture. *Human Brain Mapping*, 5:218–221, 1997.