bioRxiv preprint doi: https://doi.org/10.1101/580597; this version posted April 14, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Automated segmentation of cortical layers in BigBrain reveals divergent cortical and laminar thickness gradients in sensory and motor cortices

Konrad Wagstyl*[1,2], Stéphanie Larocque[3], Guillem Cucurull[3], Claude Lepage[1], Joseph Paul Cohen[3], Sebastian Bludau[4], Nicola Palomero-Gallagher[4,5], Thomas Funck[1], Hannah Spitzer[4], Timo Dicksheid[4], Paul C Fletcher[2], Adriana Romero[3,6], Karl Zilles[4], Katrin Amunts[4], Yoshua Bengio[3], Alan C. Evans[1]

1. McGill Centre for Integrative Neuroscience, Montreal Neurological Institute, Montreal, Canada
2. Department of Psychiatry, University of Cambridge, Cambridge, UK
3. MILA, Université de Montréal, Montreal, Canada
4. Institute of Neuroscience and Medicine (INM-1), Research Centre Jülich, 52425 Jülich Jülich Forschungszentrum, Germany
5. Department of Psychiatry, Psychotherapy and Psychosomatics, Medical Faculty, RWTH Aachen University, Aachen
6. Department of Computer Science, McGill University, Montreal, Canada

Abstract

Large-scale *in vivo* neuroimaging datasets offer new possibilities for reliable, well-powered measures of interregional structural differences and biomarkers of pathological changes in a wide variety of neurological and psychiatric diseases. However, so far studies have been structurally and functionally imprecise, being unable to relate pathological changes to specific cortical layers or neurobiological processes. We developed artificial neural networks to segment cortical and laminar surfaces in the BigBrain, a 3D histological model of the human brain. We sought to test whether previously-reported thickness gradients, as measured by MRI, in sensory and motor processing cortices, were present in a histological atlas of cortical thickness, and which cortical layers were contributing to these gradients. Identifying common gradients of cortical organisation enables us to meaningfully relate microstructural, macrostructural and functional cortical parameters.

Analysis of thickness gradients across sensory cortices, using our fully segmented six-layered model, was consistent with MRI findings, showing increasing thickness moving up the processing hierarchy. In contrast, fronto-motor cortices showed the opposite pattern with changes in thickness of layers III, V and VI being the primary drivers of these gradients. As well as identifying key differences between sensory and motor gradients, our findings show how the use of this laminar atlas offers insights that will be key to linking single-neuron morphological changes, mesoscale cortical layers and macroscale cortical thickness.

Introduction

bioRxiv preprint doi: https://doi.org/10.1101/580597; this version posted April 14, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

The cerebral isocortex has six cytoarchitectonic layers that vary depending on cortical area and local morphology. Previous work has described relationships between cytoarchitecture, cortical morphology and functional organisation [1–3]; nevertheless, the precise nature of these relationships are unclear. A core problem is that the most convenient and flexible techniques for measuring the cortex over a large population – *in vivo* MRI – lack spatial resolution while spatially precise histological approaches present profound technical problems. For example, MRI can be used to demonstrate that sensory cortices exhibit a measurable thickness gradient but cannot directly link these gradients to underlying laminar or histological patterns and are susceptible to biases from parallel gradients of myelination. Conversely, measurement problems in serial slicing of convoluted cortical tissue make thickness gradients hard to ascertain over an entire cortical region. However, creating a cortical layer segmentation of the BigBrain, a 3D histological atlas of the human brain, offers a solution to these problems and allows us to relate standard MRI measures to laminar patterns. Using this atlas, we can determine whether cortical thickness gradients can be seen in measurements made with much greater spatial resolution and whether similar cortical thickness gradients are present in fronto-motor cortices. Furthermore, it becomes possible to examine which cortical lamina contribute to these thickness gradients.

Sensory processing hierarchies describe the concept that the cerebral cortex is organised with gradually changing structural and functional properties from primary sensory areas, to secondary sensory areas and ultimately higher-order association areas. Multiple measurement modalities converge on similarly ordered patterns including single-neuronal morphological[4] and electrophysiological characteristics[5], laminar connectivity patterns of projecting cortical neurons [6,7], laminar differentiation[8], MRI cortical thickness[3], MRI myelination[9], receptor densities [10] and temporal dynamics[11]. Topographically, hierarchies are organised such that progressively higher cortical areas are located with increasing geodesic distance across the cortical surface from their primary areas [3,12]. Ordering cortical areas along these gradients provides a framework for understanding the relationships between interareal structural and functional characteristics.

Cortical thickness is a widely used marker of *in vivo* and *ex vivo* cortical structure[13]. Early histological studies noted marked interareal thickness differences on *post mortem* histological sections[14,15], which have since been replicated [16,17] and extended using *in vivo* MRI [3], and alterations in this pattern may be seen in neuropsychiatric illness [18]. MRI cortical thickness demonstrated patterns of cortical thickness relating to functional and structural hierarchical organisation across sensory cortices of both macaque and humans [3]. A similar gradient of increasing MRI cortical thickness has been reported, extending from the primary motor cortex to the anterior frontal cortex [19]. However, while classical studies of cortical histology observed that primary sensory regions are thinner than their surrounding secondary sensory cortices, the inverse is true for the motor cortex which is especially thick [14,15]. Furthermore, thickness gradients identified in MRI extended far beyond neighbouring secondary areas into association cortical areas. Thus it remains unclear whether thickness gradients found in MRI are artefactual, driven by gradient differences in cortical myelination causing systematic cortical reconstruction errors [9,20], or truly represent the underlying histology. Moreover the isocortex has six

bioRxiv preprint doi: https://doi.org/10.1101/580597; this version posted April 14, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

histologically defined layers, and it is unclear to what extent changes in cortical thickness are driven by different layers.

Carrying out analyses of histological thickness gradients poses several methodological challenges. First, most histology is carried out in 2D, with associated measurement artefacts due to oblique slicing [21]. Second, manual measurement is associated with observer dependent variability, estimated to be up to 0.5 mm [15]. Third, due to the labour-intensive nature of histological analyses, many histological atlases have a small number of sample points. These factors hinder the ability to detect and map potentially subtle, cross-cortical changes in cytoarchitecture as well as overall and laminar thicknesses. BigBrain offers a unique dataset to resolve histological cortical layers in 3D, thereby providing a concrete link between microscale patterns of structure and *in vivo* markers.

We therefore set out to automate segmentation of cortical layers in 3D in order to explore patterns of cortical and laminar thickness across sensory and fronto-motor cortical hierarchies. To do this we used a convolutional neural network to segment profiles of histological intensity sampled between the pial and white matter. Training profiles were generated from examples of manually segmented layers on cortical regions from 2D histological sections of the BigBrain. The trained network was used to segment intensity profiles derived obliquely through the 3D histological volume and generate mesh segmentations of six cortical layers. These surfaces were used to calculate cortical and laminar thicknesses. Geodesic surface distance from primary visual, auditory, somatosensory and motor areas were calculated and used as a marker of hierarchical progression. Cortical and laminar thickness gradients were calculated for each system.

bioRxiv preprint doi: https://doi.org/10.1101/580597; this version posted April 14, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Results:

Neural network training:

In the cross-validation, average per-point accuracy on the test fold was 83% +/-2% prior to post processing, suggesting the network was able to learn generalisable layer-specific features and transfer them to novel cortical areas. The predictions of the model trained on the full dataset were used to create a 3D segmentation of the cortical layers in both hemispheres (Figure 1).

Confidence results:

Layer confidence maps, given by difference between prediction values (between 0 and 1) of the highest and second highest predicted classes for each point, give an approximation of the reliability of laminar segmentations for the cortex where ground truth manual segmentations have not been carried out (Supplementary Figure 2). Throughout the cortex, the network has high confidence for supra-pial and white matter classes. Layers I, II, III, V and VI also exhibit relatively consistent confidence maps, whereas layer IV demonstrates a high variability of confidence, which is high in the occipital cortex but low anteriorly. This pattern matches with visual observations that layer IV is relatively prominent in the occipital cortex, but more difficult to identify in frontal lobe areas, and absent in primary motor cortex.

Resolution results:

Downsampling BigBrain to decrease the resolution, from 20 microns down to 1000 microns, progressively decreased the accuracy of the network on the test folds from 85% to 40% (Supplementary Figure 3). However, at 100um, the approximate upper limit for current high-resolution structural MRI, profiles had sufficient detail to maintain an accuracy of 76%.

Visual inspection of 3D layers of the cortex:

On visual inspection, the automatically identified cortical layers closely follow bands of intensity within the BigBrain (Figure 1) and continue to follow the same features beyond the limits of training examples (Figure 2a).

In the original surfaces, the white surface was placed at the maximum intensity gradient between gray matter and white matter, while the new white surface is based on the network trained on manual placement of the white boundary, which is defined by the presence of cortical neurons. On closer inspection, the maximum gradient appears to be at the border between sublayers VIa and VIb where the change in neuronal density is sharper than at the boundary between white matter and layer VI (Supplementary Figure 4).

A second feature apparent on visual inspection is that the layers do not follow a single set of rules - they vary between cortical areas. This is most readily apparent at the V1-V2 boundary where layer IV is dramatically different (Figure 2b). Layer IV is particularly thick in V1 and has multiple sublayers creating extra peaks and troughs in the intensity profiles whereas in V2 it is much thinner and no longer differentiated into sublayers. The transition from a thick layer IV to a thin layer IV occurs precisely at the boundary between these two regions suggesting the network is internally learning certain areal properties also.

bioRxiv preprint doi: https://doi.org/10.1101/580597; this version posted April 14, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Visual comparison of total and layer thickness maps

On visual inspection, maps of BigBrain cortical thickness are consistent with classical atlases of histological thickness reported by von Economo and Koskinas (Figure 3). In particular the motor cortex is the thickest part of the cortex with values over 3.8mm uncorrected (over 4.5mm when adjusted for shrinkage). The thickness of the motor cortex is often underestimated in MRI thickness measurement[22], probably due to the high degree of intracortical myelination which affects the grey-white contrast, causing the white matter surface to be too close to the grey surface, such that cortical thickness is underestimated [3,20]. Clearly visible in both BigBrain and von Economo is the thin occipital cortex.

On comparison with the available layer maps for layers I, III, V and VI from the von Economo atlas, there are several similarities but also meaningful differences (Figure 4). Layer III is thick in the precentral areas, and particularly thin in the primary visual cortex. Layers V and VI are thick in frontal and cingulate cortices, but also thin in the occipital cortex.
Of interest is the clear boundary exhibited by layer IV at the boundary between V1 and V2 in the occipital cortex. This change in thickness is abrupt enough to generate an automated anatomical label for V1 (Figures 2 and 4).

Cortical gradients and processing hierarchies

Cortical thickness was positively correlated with geodesic distance in visual (right, R=0.42, p=0; left, R=0.56, p=0), somatosensory (right, R=0.30, p=0; left, R=0.20, p=0), and auditory cortices (right, R=0.13, p=0; left, R=0.20, p=0) (Figure 5A-C). By contrast in the motor cortex, thickness was negatively correlated with geodesic distance (right, R=-0.16, p=0; left, R=-0.25, p=0) (Figure 5D). This is consistent with MRI thickness findings in sensory gradients but contradictory for the fronto-motor gradient.

Cortical layers did not contribute equally to the total thickness gradient in the visual cortex and somatosensory (Figure 6A). Layers III and V had the largest contributions to the total thickness gradient, followed by layer VI and then II. A similar but less pronounced pattern was seen within the auditory cortex.

In the motor cortex the inverse was true, with decreases in layers III, V and VI. Layer IV also exhibited a decrease in thickness, however there was low network confidence for layer IV in these areas and thus this particular finding is inconclusive.

Thus, consistent with MRI thickness gradients, somatosensory and visual areas exhibited positive histological thickness gradients primarily driven by layers III, V and VI. By contrast the fronto-motor areas exhibited an inverse gradient, peaking in the motor cortex and driven by the same layers (Figure 6B & C).

Open data availability

All code used for the segmentation and analysis of cortical layers will be made available at https://github.com/kwagstyl and BigBrain volumes, layer surfaces and intensity profiles are available to download at ftp://bigbrain.loris.ca/.

bioRxiv preprint doi: https://doi.org/10.1101/580597; this version posted April 14, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Discussion

We automatically segmented the six histological layers of the cerebral cortex in the BigBrain and used these to test for gradients of cortical and laminar thickness within sensory and motor processing hierarchies. Consistent with previous findings using *in vivo* MRI and 2D histological measurements [15], sensory hierarchies exhibited a gradient of increasing cortical thickness from primary sensory to higher order areas. In visual, somatosensory and auditory cortices, these gradients were primarily driven by layers III, V, VI. By contrast, the fronto-motor cortices exhibited a decreasing cortical thickness gradient away from the primary motor cortex towards higher order frontal areas, which was driven by decreases in these layers. These findings highlight the utility of the BigBrain for linking micro- and macro-scale patterns of cortical organisation.

Gradients of thickness are large scale markers of systematic changes in the cortical microcircuit. The volume of the cortex is 80-90% neuropil [23,24], of which 60% is axons and dendrites and the remainder is synaptic boutons, spines and glia. As neuronal density decreases with increasing cortical thickness [25,26], and most of the volume of the cortex is neuropil, increased thickness is most likely to mark increased intracortical connectivity [13]. At a laminar level, the strongest contributors to the overall thickness gradients were layers III, V and VI (Figure 6). Cell-morphological studies in macaques have shown that the cell size and dendritic arborisation of layer III and V pyramidal neurons increase across the visual pathway [4,27,28]. Similarly, afferent axonal patch sizes scale with pyramidal neuronal arborisation [29]. Increasing dendritic arborisation, axonal field size and number of synapses all contribute to an increase in the volume of laminar neuropil and are therefore a likely driver of the laminar and overall thickness gradients measured here. Layer VI also increased in thickness in sensory pathways. However, while neurons located in these layers might exhibit increases of their associated neuropil, the measured thickness change may in part be due to the extended arborisations of III/V pyramidal neurons forming connections within these layers. Histological gradients of layer thickness provide us with a meso-scale link between *in vivo* patterns of MRI cortical thickness and microstructural, neuronal-scale changes in the cortical microcircuit. Such links help us to understand the neurobiological significance of interindividual, longitudinal and neuropathological biomarkers[13].

In contrast to *in vivo* studies of fronto-motor functional, myelin and MRI cortical thickness organisation, which place the primary motor cortex at the same level as primary sensory areas[9,12,19], we found that total and laminar fronto-motor thickness gradients were the inverse of those measured in sensory cortices. Histologically, the motor cortex was especially thick and the thickness decreased with geodesic distance from the primary motor cortex, with layers III, V and VI following a similar inverse pattern. This finding is consistent with reported trends in other histological properties such as laminar structural type [15] and neuronal density [25], as well as the observation that the motor cortex has large, pyramidal neurons with extensive dendritic arborisation [30,31]. Functionally, these structural differences might be considered in terms of narrow, specific columnar receptive fields for accurate sensory perception [32], and wider receptive fields [33] for the coordination of multiple muscle groups [34] in precise motor control. Thus

bioRxiv preprint doi: https://doi.org/10.1101/580597; this version posted April 14, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

there is a coherent group of cortical microcircuit properties which diverge from patterns of cortical myelination and fMRI-derived gradients, locating the motor cortex at the peak of a gradient of increasing cortical thickness, layer III, V and VI thickness and pyramidal neuronal arborisation, with primary sensory areas at the opposite extreme.

Atlas of cortical layers

The layers we have generated to test gradient based hypotheses have applications beyond the scope of this study. Surface-based models of layer structure also create a framework for translating between microstructural modalities and surface-based neuroimaging. Furthermore, current approaches to measuring laminar structure and function *in vivo* rely on prior models of the cortical layers - for example signal-source simulation in MEG [35] or for sampling laminar BOLD signal in fMRI [36]. The whole-brain histological models for areal layer depth provided here, combined with a good understanding of how the layers vary with local cortical morphology[21,37,38] will aid such anatomical models.

Limitations

It is important to acknowledge that the gradients of laminar thickness measured may be affected by limitations in the BigBrain dataset. The first limitation is that the *post mortem* brain was damaged during extraction and mounting. In some areas this resulted in minor shears. This problem was addressed to some extent through the utilisation of non-linear registration techniques. Nevertheless, some shifts in cortical tissue between consecutive sections are present and will affect the accuracy of layer reconstructions. In other areas, the cortex has been torn. Spatial smoothing and the large total number of sample points make it unlikely that these errors are affecting the results. A second limitation is that there is only one BigBrain. Future work will be necessary to establish the interindividual and age-dependent variability in laminar structure, either using other histological BigBrains or with complementary high-resolution MRI imaging approaches.

Total cortical thickness and thicknesses for each of the six isocortical layers were measured in the BigBrain to explore the histological drivers of MRI-based thickness gradients. Overall, the pattern of thickness in the BigBrain was consistent with histological atlases of cortical thickness. In the visual and somatosensory cortices, an increasing gradient of histological cortical thickness was identified and found to be primarily driven by layers III, V and VI. In the fronto-motor cortex, the inverse pattern was found. These findings provide a link between patterns of microstructural change and morphology measurable through MRI and emphasise the importance of testing MRI-based anatomical findings against histological techniques. The laminar atlases provide an invaluable tool for comparison of histological and macroscale patterns of cortical organisation.

Acknowledgements

Parts of this work have received support from Healthy Brains for Healthy Lives, the Avrith MNI-Cambridge Neuroscience Collaboration award, and the European Union's Horizon

bioRxiv preprint doi: https://doi.org/10.1101/580597; this version posted April 14, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

2020 Framework Programme for Research and Innovation under Grant Agreement No 720270 (Human Brain Project SGA1).

Author Contributions

Conception and design of the work: KW, KA, YB, ACE. Acquisition: CL, SB NPG, KZ, KA. Analysis and interpretation of data: KW, SL, GC, JPC, AR, KZ. Creation of new software used in the work: KW, SL, GC,CL, JPC, AR, TF, HS, TD. Drafted the work or substantively revised it: KW, SL, JPC, CL, NPG, ACE.

Competing Interests statement

The authors have no competing interests relating to this work.

bioRxiv preprint doi: https://doi.org/10.1101/580597; this version posted April 14, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Methods:

Volumetric data preparation:

BigBrain is a 20x20x20μm (henceforth described as 20μm) resolution volumetric reconstruction of a histologically processed *post mortem* human brain (male, aged 65), (available for download at https://bigbrain.loris.ca). In order to run computations on this amount of data, the BigBrain was partitioned into 125 individual blocks, corresponding to five subdivisions in the x, y and z directions, with overlap. The overlap of blocks was calculated to be sufficient such that each single cortical column could be located in a single block, enabling extraction of complete intensity profiles between pairs of vertices at the edge of blocks without intensity values being altered by boundary effects when the data were smoothed. Blocks were smoothed anisotropically[39], predominantly along the direction tangential to the cortical surface, to maximise interlaminar intensity differences while minimising the effects of intralaminar intensity variations caused by artefacts, blood vessels and individual neuronal arrangement[21]. The degree of anisotropic smoothing is determined by repeatedly applying the diffusive smoothing algorithm. The optimal level of smoothing was previously determined and gave an effective maximum full width at half maximum (FWHM) of 0.163mm[21]. For subsequent analyses both the raw 20μm and anisotropically smoothed blocks were used.

Lower resolution volumes were extracted by subsampling the raw BigBrain 20μm volume at 40, 100, 200, 400 and 1000μm. Anisotropically smoothed volumes were also generated at each of these resolutions.

Profile extraction:

Pial and white surfaces originally extracted using a tissue classification of 200μm were taken as starting surfaces[40]. Prior to intensity profile extraction, in order to minimise cortical streamlines intersecting, to improve their approximation of columnar trajectories and to ensure profiles captured the full extent of the cortex, the following steps were taken:

1) A mid-surface was generated that was closer to the pial surface in sulci and closer to the white surface in gyri, weighting the distance vector by the cortical curvature. Thus the mid surface was closer to the surface with the higher curvature.
2) This mid-surface was upsampled to 655362 vertices to increase the resolution of the surface.
3) For each vertex, the vector between nearest points on the pial and white surfaces was calculated.
4) To avoid crossing profiles which can result in mesh self-intersections, the vector components were smoothed across the mid-surface with a FWHM of 3mm (Supplementary Figure 1).
5) Profiles were calculated along these vectors from the midsurface, extending the profiles 0.5mm further than the minimum distance in the pial or white direction, to ensure the resultant intensity profile captured the full extent of the cortex.

The resultant profiles were less oblique and more likely lined up with the cortical columns (Supplementary Figure 1).

bioRxiv preprint doi: https://doi.org/10.1101/580597; this version posted April 14, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Extended intensity profiles were then created by sampling voxels at 200 equidistant points between each pair of vertices from the raw and anisotropically smoothed BigBrain volumes, at each available resolution. For an extended profile of approximately 4mm, this gives a distance of 0.02mm or 20µm between points, corresponding to the highest resolution volume available. Profile intensity values were adjusted by regressing between mean profile intensity and posterior-anterior coordinate in 3D space, to account for the decreasing staining intensity.

Training data:

Manual segmentations of the 6 cortical layers were created on 51 regions of the cortex, distributed across 13 of the original histological BigBrain sections at 5µm (Figure 2). Layer I, the molecular layer, is relatively cell-sparse with few neurons and glia. Layer II, the external granular layer, is a much more dense band of small granular cells. Layer III, the external pyramidal layer, is characterised by large pyramidal neurons that become more densely packed towards its lower extent. Layer IV, the internal granular layer (usually referred to simply as the “granular layer”), generally contains only granular neurons, bounded at its lower extent by pyramidal neurons of layer V. Layer V, the internal pyramidal layer, contains large but relatively sparse pyramidal neurons while layer VI, the multiform layer, has a lower density of pyramidal neurons [14]. Segmentations were inspected by histological experts: SB, NPG and KZ. This resolution is sufficient to distinguish cell body distribution patterns used to delineate specific layers. Averaged across all training examples, layer classes contributed to profiles as follows: background/csf: 14.6%, layer I: 7.5%, layer II: 5.6% , layer III: 20.8%, layer IV: 5.5%, layer V: 14.8%, layer VI: 17.8%, white matter: 13.4%. For the cortical layers, these values represent an approximate relative thickness.

Manual segmentations were then co-registered to the full aligned BigBrain 3D space. The manually drawn layers were used to create corresponding pial and white surfaces and these cortical limits were extended into the white matter and beyond the pial surface between 0.25mm and 0.75mm so as to match the variability of cortical extent in the test profile dataset. Training profiles were created sampling raw, smoothed and manually segmented data, generating thousands of profiles per sample. Each pixel in the labelled data had a class value of 0-7, where pixels superficial to the pial surface were set to 0, followed by layers numbered 1-6, and white matter was classed as 7. This 1D profile based approach greatly expanded the training dataset from 51 labelled 2D samples to 564,041 profiles.

Neural network:

A 1D convolutional network for image segmentation was created with the following structure. The network was created using stacked identical blocks. Each block contained a batch normalisation layer to normalise feature distributions between training batches, a rectify non-linearity layer used to model neurons which can have a graded positive activation but no negative response[41], and a convolutional layer[42]. There was a final convolutional layer with filter size 1 and 8 feature maps, one for each class. The cost function was median class-frequency weighted cross-entropy. Class-frequency weighting was added to weigh errors according to the thickness of the layers so that incorrectly classified points in thinner layers were more heavily weighted than errors in incorrectly classified thicker layers[43]. Raw and smoothed profiles were

bioRxiv preprint doi: https://doi.org/10.1101/580597; this version posted April 14, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

considered as two input channels. The network was iteratively trained until the accuracy did not improve for 50 epochs (all training profiles are seen once per epoch). At this point, the previous best model was saved and used for subsequent testing on the full dataset. When testing the network, a soft maximum was then applied to detect the most likely layer class for each point. The output was a matrix of 8 (predicted layers) by 200 (sample points) by 655362 (vertices on a mesh) by 2 (cortical hemispheres).

For each vertex, a measure of confidence was calculated from these predictions. Per point confidence is the difference between the prediction value for the highest predicted class and the value of the 2nd highest predicted class. Per class/layer confidence is the mean confidence for all points in that class/layer. The per-vertex summary measure is the mean across all points in the profile. These measures give an indication of the relative confidence for the regional and laminar classifications.

Hyperparameter optimisation and cross-validation:

There is no consensus method for finding optimum parameters for a neural network. Here, a set of 50 experiments with random hyperparameters was carried out to explore their impact on training accuracy. Learning rate, convolutional filter size, number of layers (blocks), weight decay and batch size were all varied. In summary, the final network was initialised with 6 layers, filter size = 49, learning rate = 0.0005, weight decay = 0.001, where the learning rate determines the amount weights are updated on each iteration and weight decay determines the rate at which weights decrease each iteration, which helps prevent overfitting.

For network cross validation, the manually labelled areas were subdivided into 10 equally sized random subsets or folds. Initially, 2 folds were removed from the dataset during training and network weights were optimised for segmenting samples on one of these folds. This trained network was then used to predict layers on the final previously unseen test fold from which the accuracy was calculated. This process was repeated 10 times to generate an estimate of the network's ability to segment novel cortical regions. The same process was carried out using profiles extracted at all available resolutions.

For generating BigBrain layer segmentations, the network was trained on the full training dataset and tested on all intensity profiles.

Shrinkage estimate:

Histological processing, including fixation and sectioning, causes distortion of the tissue that is non-uniform in the x,y and z directions. Part of this distortion was corrected in the original reconstruction of the BigBrain [44].

Initial shrinkage of the brain during fixation prior to sectioning was calculated based on the estimated fresh volume of BigBrain, based on the original fresh weight, and the volume after histological processing. This gave a volume-based (3D) shrinkage factor of 1.931, which corresponds to an isotropic length-based (1D) shrinkage factor of 1.245.

To estimate the extent of shrinkage in each of the three orthogonal directions, the BigBrain volume was linearly coregistered to a MRI volumetric template based on a group of older subjects (ADNI) [45]. The transformation matrix gave linear scale factors of 1.15, 1.22 and 1.43 in the x,y and z directions, with a mean of 1.26. The concordance of these measures of

bioRxiv preprint doi: https://doi.org/10.1101/580597; this version posted April 14, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

shrinkage suggests that subsequent thickness and length estimates can be adequately corrected for comparison to *in vivo* measures.

Thus, to approximately compensate for the non-uniform compression of xyz, we transformed the mesh surfaces into MNI space based on the ADNI template. Subsequent thickness analyses were carried out on the transformed meshes. Such compensation for shrinkage is necessary when analysing cortical thickness gradients on oblique profiles in 3D over the whole brain.

Surface reconstruction: post-processing 1D profiles:

1D classified profiles were transformed into mesh layer boundary reconstructions as follows. Transitions between predicted layers were located for each profile and the coordinates of these transitions became vertex locations for the new layer meshes. For the small number of vertices where the network failed (less than 1%), vertex locations were interpolated from the neighbouring vertices. Surface indices were smoothed 0.5mm FWHM across the cortical surface and 20 iterations of shrinkage-free mesh smoothing was applied to the output surface [46]. This removed non-biologically high frequency changes in surface curvature, most commonly due to minor, local misalignment of consecutive 2D coronal sections.

Cortical thickness, layer thickness:

Cortical thickness was calculated between pial and white cortical surfaces and laminar thicknesses were calculated between adjacent pairs of cortical surfaces.

Masking:

Manual masks were created to remove medial wall, archicortex, including areas of the cingulate and entorhinal cortex which do not have 6 layers, and large cuts in the anterior temporal cortex (caused by the saw during extraction of the brain from the skull) from subsequent analyses.

Gradients and processing hierarchies:

Surface labels for the primary visual, somatosensory, auditory and motor areas were manually delineated on each hemisphere using morphological markers and histological characteristics (Supplementary Figure 5a). Geodesic distance across the cortex from each of these labels was calculated to each of the vertices within a larger masked area containing associated cortical regions (Supplementary Figure 5b&c) [3,19,47,48].

bioRxiv preprint doi: https://doi.org/10.1101/580597; this version posted April 14, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Figures:

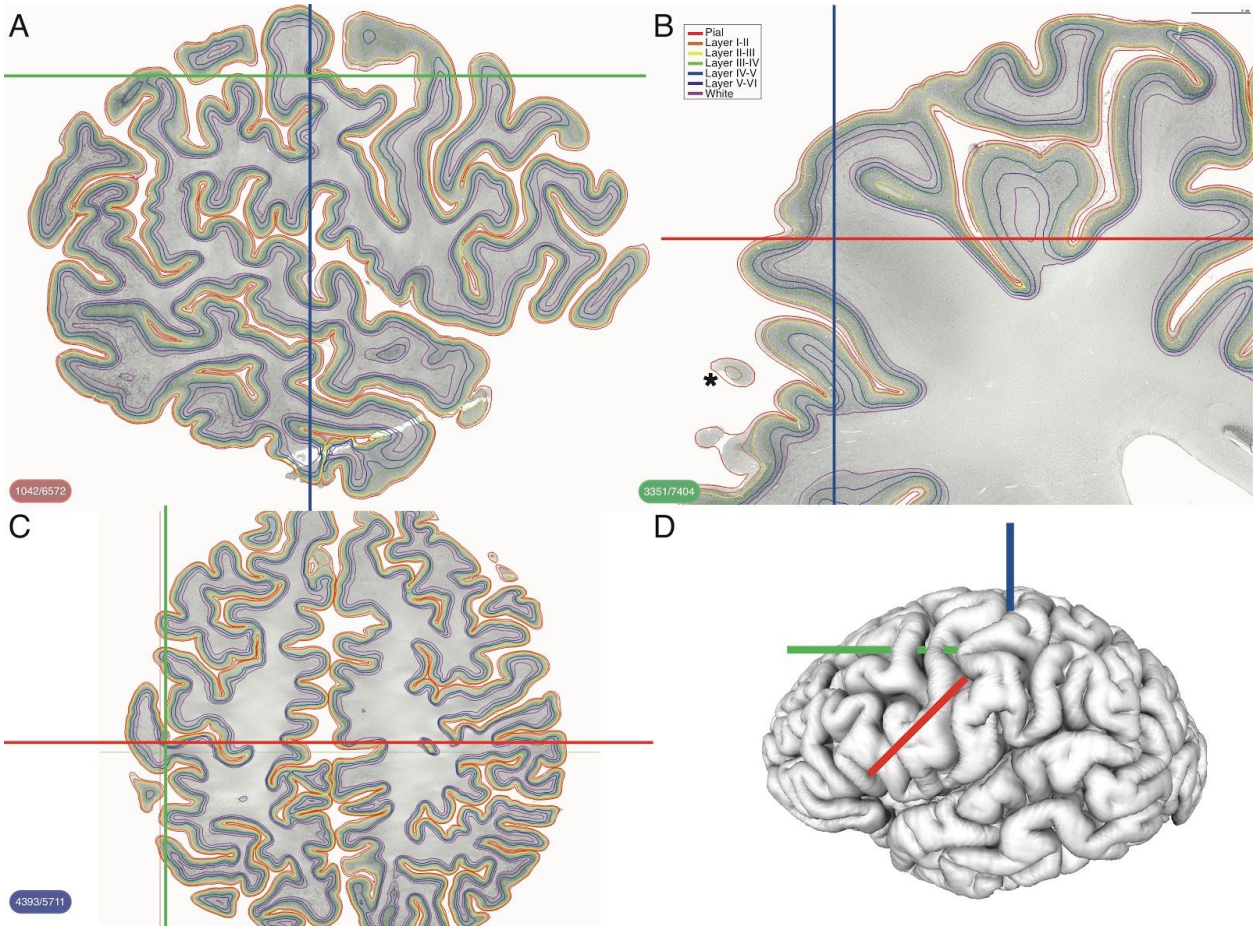

Figure 1

Cortical layers in 3D. Six cortical layers segmented on the 3D volume on three orthogonal planes: A=sagittal, B=coronal, C=axial. Panel D shows the location of the sections on the reconstructed pial surfaces of the 3D BigBrain. B, the coronal plane is the original plane of sectioning. Within this plane is marked an area of the cortex where layers would be difficult to segment in 2D due to the oblique sectioning of the cortex.

bioRxiv preprint doi: https://doi.org/10.1101/580597; this version posted April 14, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

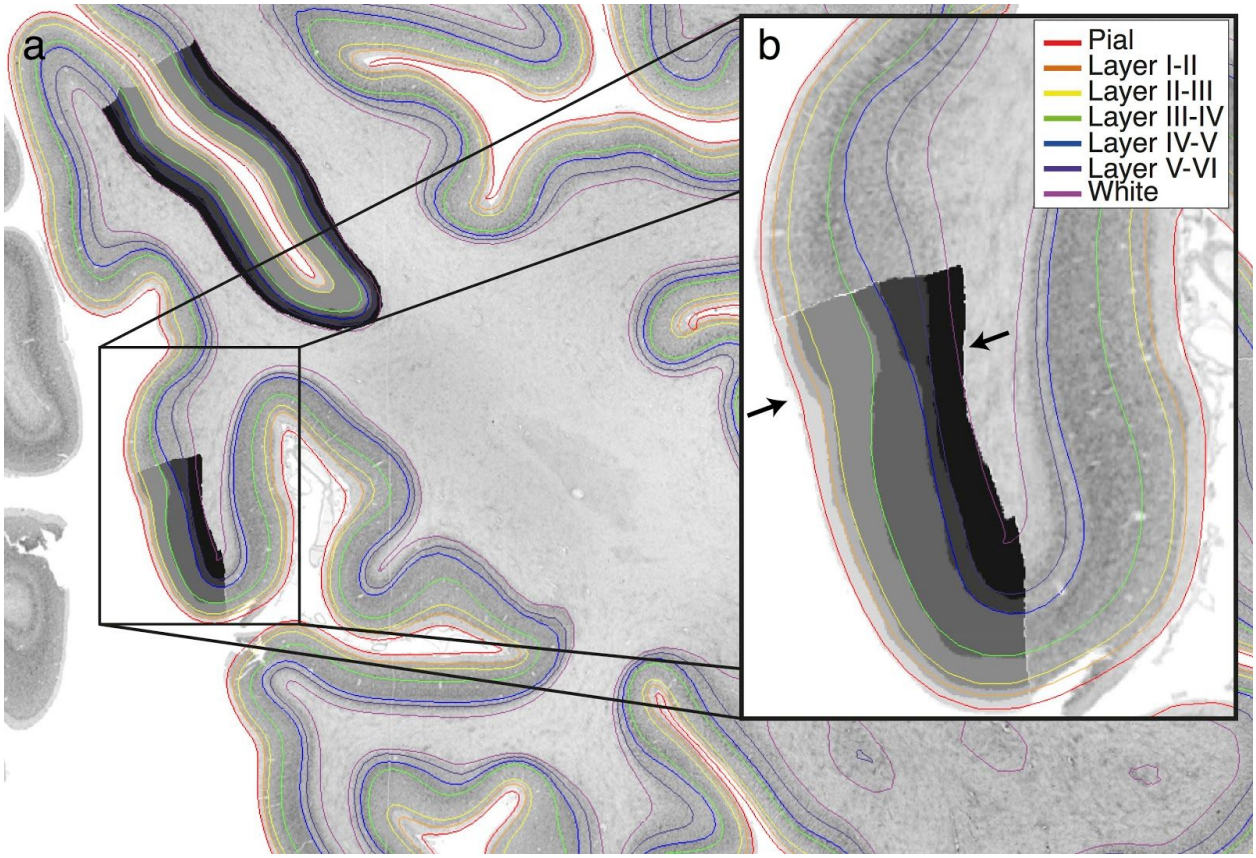

Figure 2

Cortical layers intersected on a 2D section with manually segmented layers. a) The boundaries clearly follow the same contours delineated by the manually segmented training areas, and appear to accurately follow the layer bounds outside of each training area. b) At the V1-V2 boundary (marked with arrows) the thickness of layer IV changes dramatically in both manual and automated segmentations (between red and yellow lines), with additional peaks in V1 intensity due to the sublayers of layer IV. As each profile is individually segmented by the network, without reference to the neighbouring profiles, the network is able to apply area-specific rules according to the shape of the profile, suggesting it might be internally identifying the area from which the profile is extracted as being either V1 or V2.

bioRxiv preprint doi: https://doi.org/10.1101/580597; this version posted April 14, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

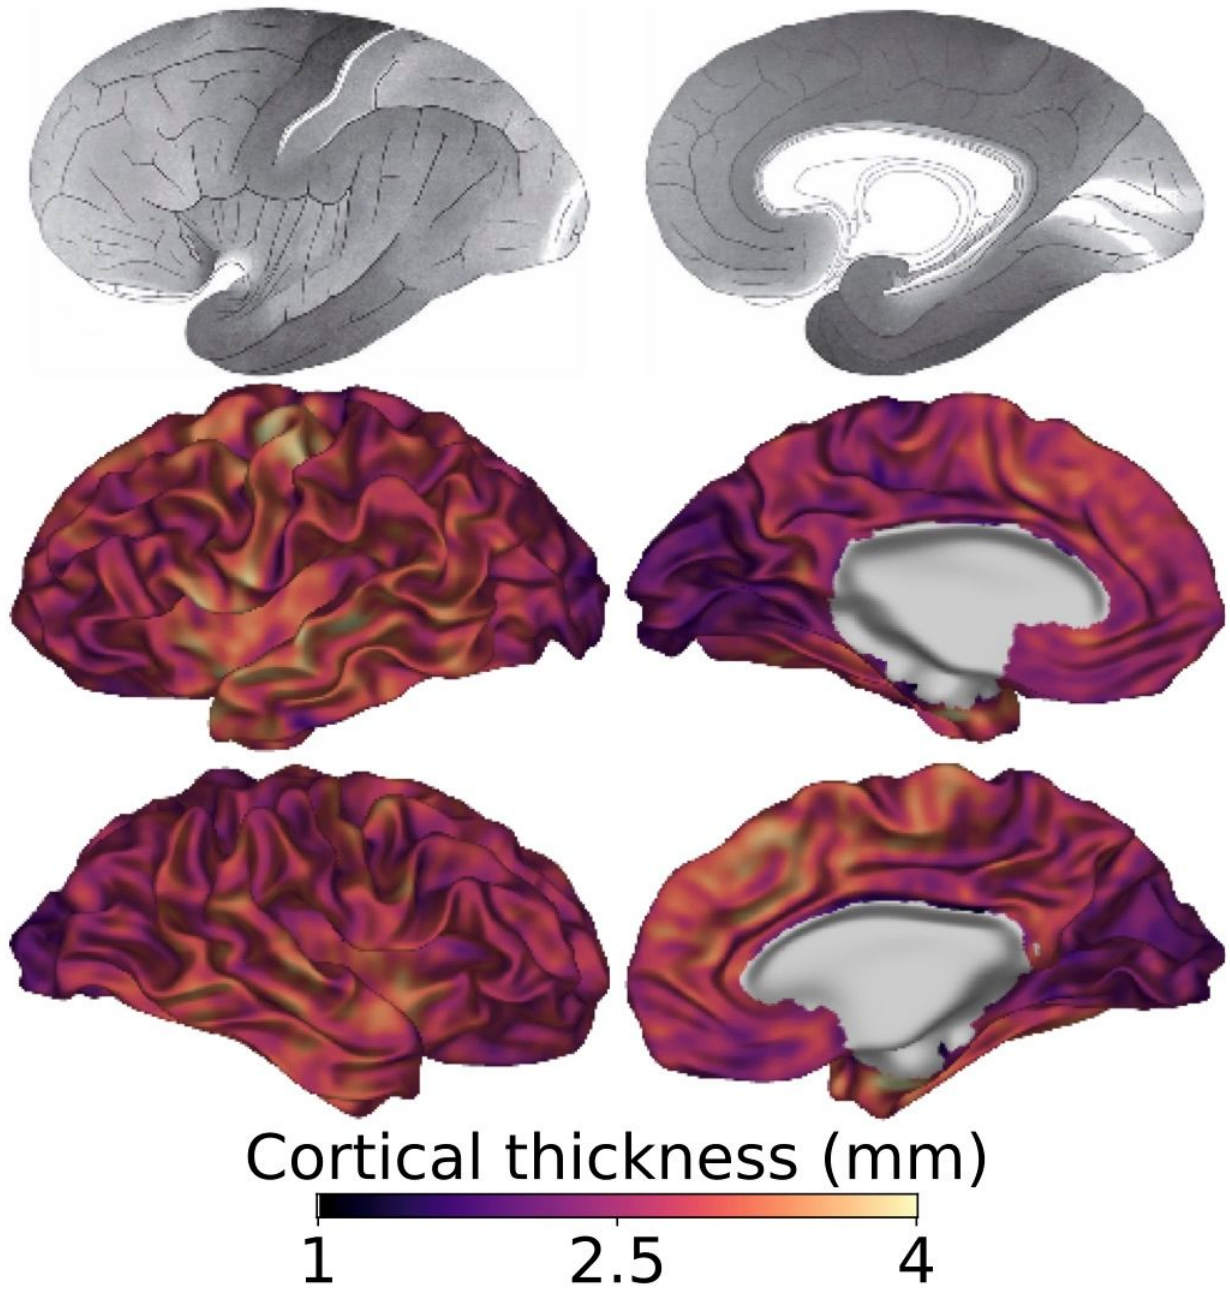

Figure 3

Comparison of von Economo's maps of cortical thickness with cortical thickness on the left and right hemispheres of BigBrain. On the von Economo maps, white is thinner cortex and black is thicker. Thickness values range from 1.8mm in the calcarine sulcus to 4.5 mm in the precentral gyrus. The pattern of cortical thickness across the BigBrain (displayed on smoothed surfaces, values were smoothed 3mm FWHM) matched that measured by von Economo and Koskinas. In particular, the precentral gyrus, or primary motor cortex, which is often underestimated with MRI, was the thickest area. Additionally, the occipital cortex around the calcarine sulcus was particularly thin in both BigBrain and von Economo. While broad trends are

bioRxiv preprint doi: https://doi.org/10.1101/580597; this version posted April 14, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

consistent between the two maps, BigBrain exhibits more local variability, for instance due to cortical folding, due to the higher density of measurements made. Von Economo reported around 60 regional thickness measurements whereas around 1 million vertices have been measured on BigBrain.

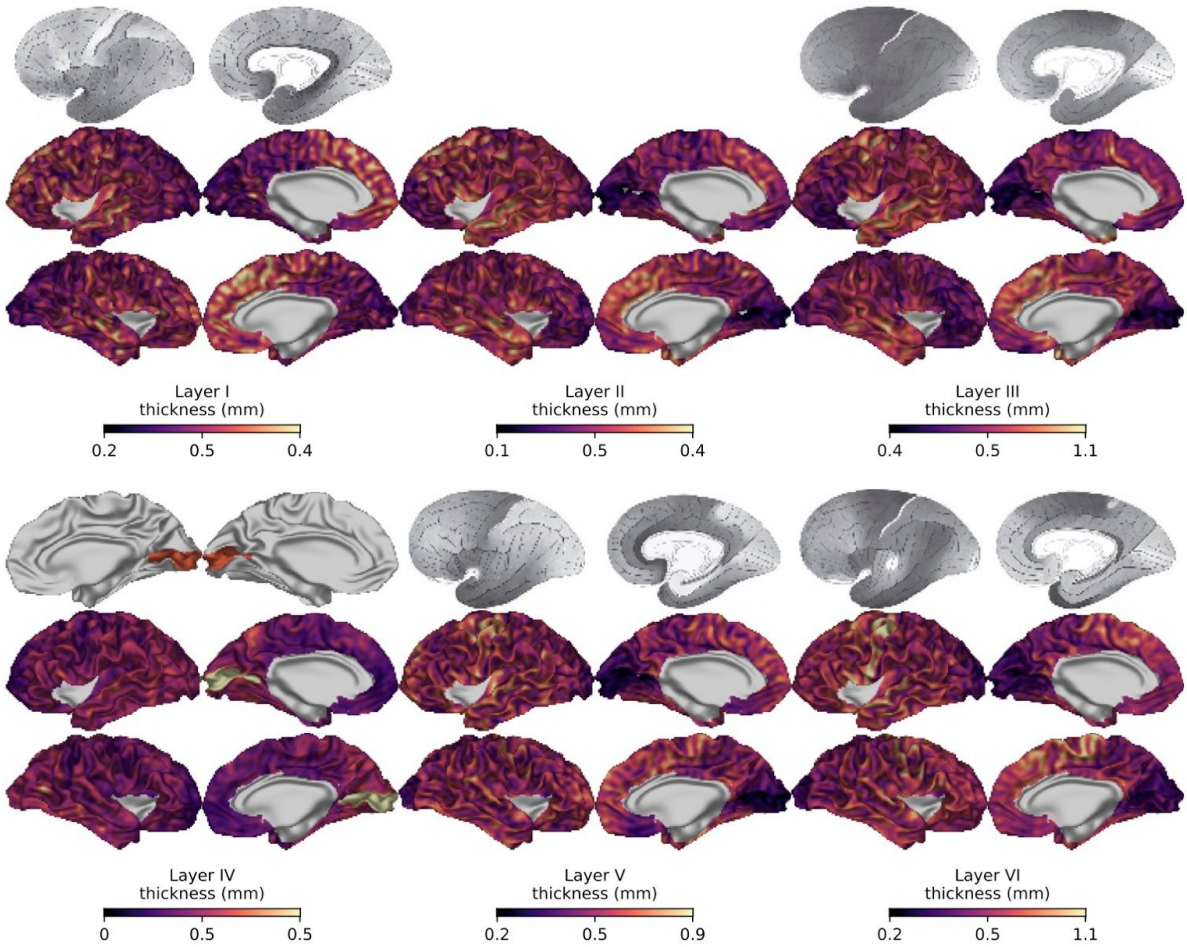

Figure 4

Comparison of von Economo's maps laminar thickness with laminar thicknesses on the left and right hemispheres of BigBrain. BigBrain were smoothed across the surface with a 3mm FWHM Gaussian kernel. From von Economo, surface mappings of laminar thickness were available for layers I, III V and VI. For these maps, white represents thinner cortex and black thicker. Similarities include the clear changes in thickness in pre and post central thicknesses of layers III, V and VI. For layer IV, the most striking feature is the abrupt change in layer IV thickness at the V1 V2 border. For reference, the manual delineation of the primary visual cortex (V1) based on cytoarchitectonic features is included. This abrupt change and the unique features of layer IV in V1 lead us to conclude that the neural network may have internally learned to recognise V1 and apply the appropriate laminar segmentation rules.

bioRxiv preprint doi: https://doi.org/10.1101/580597; this version posted April 14, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

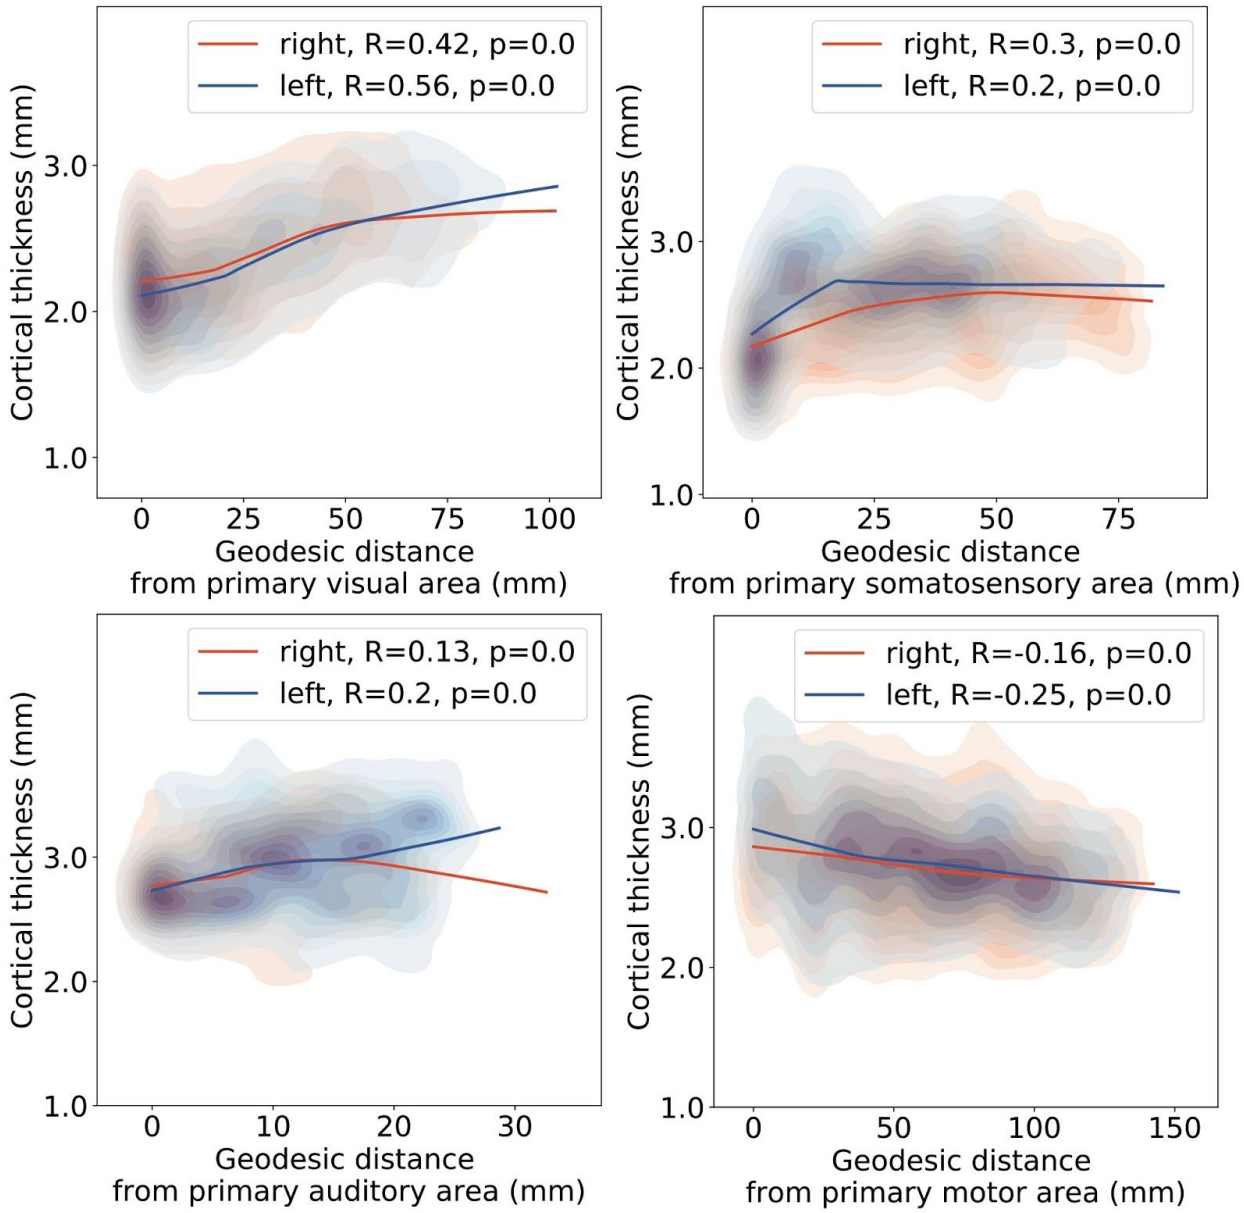

Figure 5

Cortical thickness with increasing distance from the primary area. For primary visual and somatosensory cortices, consistent with MRI studies of cortical thickness, thickness increased with geodesic distance from the primary sensory areas. In the auditory cortex, a weak relationship was present. For the motor cortex, a negative relationship was present with thickness decreasing from the primary motor cortex into the frontal cortex. This structural gradient is the inverse of the pattern of myelination and of previously reported MRI frontal thickness gradients, but consistent with patterns of structural type and neuronal density. These finding suggest the presence of distinct but overlapping structural hierarchies.

bioRxiv preprint doi: https://doi.org/10.1101/580597; this version posted April 14, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

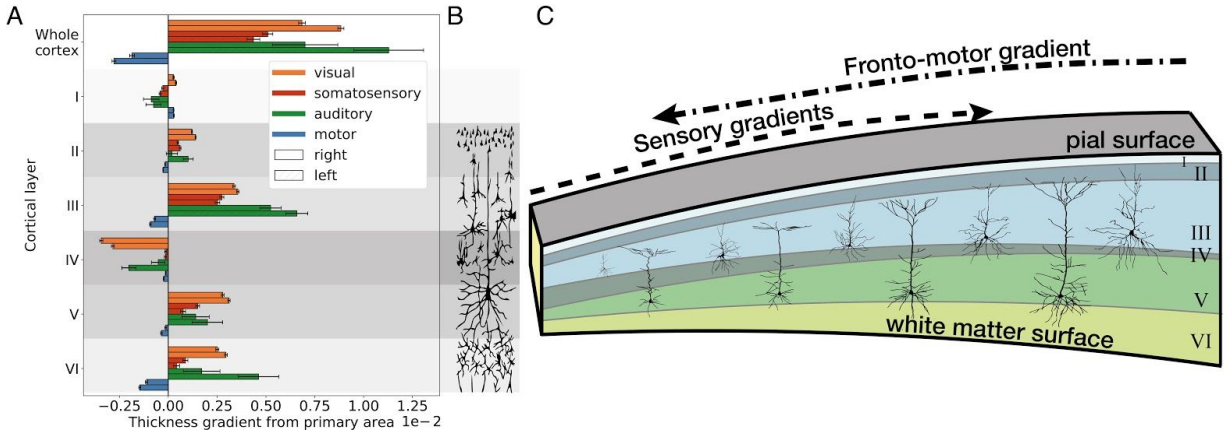

Figure 6

Gradients of cortical and laminar thickness against geodesic distance from primary areas. A) Fronto-motor gradients show an inverse relationship from sensory gradients both on cortical thickness and laminar thicknesses. Increasing sensory cortical thickness gradients were generally driven by thickness increases in layers III, V and VI. By contrast decreasing fronto-motor cortical thickness gradients exhibited decreases in thickness of the same layers. B) Typical neuronal types and morphologies of individual cortical layers. C) Cortical thickness gradients in either direction are primarily driven by changes in pyramidal cell layers. Single-cell morphological studies of these neurons in macaque sensory processing pathways reveal increasing dendritic arborisation [49] consistent with the hypothesis that laminar volume changes and ultimately thickness changes represent increases in intracortical connectivity.

bioRxiv preprint doi: https://doi.org/10.1101/580597; this version posted April 14, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

References

1. Broca, P. Remarks on the seat of the faculty of articulated language, following an observation of aphemia (loss of speech). *Bulletin de la Société Anatomique* **6**, 330–357 (1861).
2. Fischl, B. *et al.* Cortical folding patterns and predicting cytoarchitecture. *Cereb. Cortex* **18**, 1973–1980 (2008).
3. Wagstyl, K., Ronan, L., Goodyer, I. M. & Fletcher, P. C. Cortical thickness gradients in structural hierarchies. *Neuroimage* **111**, 241–250 (2015).
4. Elston, G. N. & Rosa, M. G. Morphological variation of layer III pyramidal neurones in the occipitotemporal pathway of the macaque monkey visual cortex. *Cereb. Cortex* **8**, 278–294 (1998).
5. Elston, G. N. & Fujita, I. Pyramidal cell development: postnatal spinogenesis, dendritic growth, axon growth, and electrophysiology. *Front. Neuroanat.* **8**, 78 (2014).
6. Felleman, D. J. & Van Essen, D. C. Distributed hierarchical processing in the primate cerebral cortex. *Cereb. Cortex* **1**, 1–47 (1991).
7. Goulas, A., Zilles, K. & Hilgetag, C. C. Cortical Gradients and Laminar Projections in Mammals. *Trends Neurosci.* (2018). doi:10.1016/j.tins.2018.06.003
8. Barbas, H. Pattern in the laminar origin of corticocortical connections. *J. Comp. Neurol.* **252**, 415–422 (1986).
9. Huntenburg, J. M. *et al.* A Systematic Relationship Between Functional Connectivity and Intracortical Myelin in the Human Cerebral Cortex. *Cereb. Cortex* **27**, 981–997 (2017).
10. Palomero-Gallagher, N. & Zilles, K. Cortical layers: Cyto-, myelo-, receptor- and synaptic architecture in human cortical areas. *Neuroimage* (2017). doi:10.1016/j.neuroimage.2017.08.035

bioRxiv preprint doi: https://doi.org/10.1101/580597; this version posted April 14, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

11. Murray, J. D. *et al.* A hierarchy of intrinsic timescales across primate cortex. *Nat. Neurosci.* **17**, 1661–1663 (2014).

12. Margulies, D. S. *et al.* Situating the default-mode network along a principal gradient of macroscale cortical organization. *Proceedings of the National Academy of Sciences* (2016). doi:10.1073/pnas.1608282113

13. Wagstyl, K. & Lerch, J. P. Cortical Thickness. in *Brain Morphometry* (eds. Spalletta, G., Piras, F. & Gili, T.) 35–49 (Springer New York, 2018).

14. Brodmann, K. *Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues*. (Barth, 1909).

15. von Economo, C. F. & Koskinas, G. N. *Die cytoarchitektonik der hirnrinde des erwachsenen menschen*. (J. Springer, 1925).

16. Lerch, J. P. & Evans, A. C. Cortical thickness analysis examined through power analysis and a population simulation. *Neuroimage* **24**, 163–173 (2005).

17. Fischl, B. & Dale, A. M. Measuring the thickness of the human cerebral cortex from magnetic resonance images. *Proceedings of the National Academy of Sciences* **97**, 11050–11055 (2000).

18. Wagstyl, K. *et al.* Multiple markers of cortical morphology reveal evidence of supragranular thinning in schizophrenia. *Transl. Psychiatry* **6**, e780 (2016).

19. Thiebaut de Schotten, M. *et al.* Rostro-caudal Architecture of the Frontal Lobes in Humans. *Cereb. Cortex* (2016). doi:10.1093/cercor/bhw215

20. Glasser, M. F. & Van Essen, D. C. Mapping human cortical areas in vivo based on myelin content as revealed by T1- and T2-weighted MRI. *J. Neurosci.* **31**, 11597–11616 (2011).

21. Wagstyl, K. *et al.* Mapping Cortical Laminar Structure in the 3D BigBrain. *Cereb. Cortex* **28**, 2551–2562 (2018).

bioRxiv preprint doi: https://doi.org/10.1101/580597; this version posted April 14, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

22. Zilles, K. & Amunts, K. Anatomical Basis for Functional Specialization. in *fMRI: From Nuclear Spins to Brain Functions* (eds. Uludag, K., Ugurbil, K. & Berliner, L.) 27–66 (Springer US, 2015).

23. Braitenberg, V. & Priv Doz Dr. *Anatomy of the Cortex: Statistics and Geometry*. (Springer Berlin Heidelberg, 1991).

24. Chklovskii, D. B., Schikorski, T. & Stevens, C. F. Wiring optimization in cortical circuits. *Neuron* **34**, 341–347 (2002).

25. Collins, C. E., Airey, D. C., Young, N. A., Leitch, D. B. & Kaas, J. H. Neuron densities vary across and within cortical areas in primates. *Proceedings of the National Academy of Sciences* **107**, 15927–15932 (2010).

26. la Fougère, C. *et al.* Where in-vivo imaging meets cytoarchitectonics: the relationship between cortical thickness and neuronal density measured with high-resolution [18F]flumazenil-PET. *Neuroimage* **56**, 951–960 (2011).

27. Elston, G. N. & Rosa, M. G. Pyramidal cells, patches, and cortical columns: a comparative study of infragranular neurons in TEO, TE, and the superior temporal polysensory area of the macaque monkey. *J. Neurosci.* **20**, RC117 (2000).

28. Scholtens, L. H., Schmidt, R., de Reus, M. A. & van den Heuvel, M. P. Linking macroscale graph analytical organization to microscale neuroarchitectonics in the macaque connectome. *J. Neurosci.* **34**, 12192–12205 (2014).

29. Lund, J. S., Yoshioka, T. & Levitt, J. B. Comparison of intrinsic connectivity in different areas of macaque monkey cerebral cortex. *Cereb. Cortex* **3**, 148–162 (1993).

30. Cajal, S. R. y. Estudios sobre la corteza cerebral humana. II. Estructura de la corteza motriz del hombre y mamíferos superiores. *Revista Trimestral Micrográfica* **4**, 117–200 (1899).

bioRxiv preprint doi: https://doi.org/10.1101/580597; this version posted April 14, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

31. Barbas, H. & García-Cabezas, M. Á. Motor cortex layer 4: less is more. *Trends Neurosci.* **38**, 259–261 (2015).

32. Song, C., Schwarzkopf, D. S., Kanai, R. & Rees, G. Neural population tuning links visual cortical anatomy to human visual perception. *Neuron* **85**, 641–656 (2015).

33. Schellekens, W., Petridou, N. & Ramsey, N. F. Detailed somatotopy in primary motor and somatosensory cortex revealed by Gaussian population receptive fields. *Neuroimage* **179**, 337–347 (2018).

34. Mountcastle, V. B. The columnar organization of the neocortex. *Brain* **120 (Pt 4)**, 701–722 (1997).

35. Troebinger, L., López, J. D., Lutti, A., Bestmann, S. & Barnes, G. Discrimination of cortical laminae using MEG. *Neuroimage* **102 Pt 2**, 885–893 (2014).

36. Muckli, L. *et al.* Contextual Feedback to Superficial Layers of V1. *Curr. Biol.* **25**, 2690–2695 (2015).

37. Waehnert, M. D. *et al.* Anatomically motivated modeling of cortical laminae. *Neuroimage* **93 Pt 2**, 210–220 (2014).

38. Wagstyl, K., Paquola, C., Bethlehem, R., Evans, A. C. & Huth, A. *Equivolumetric layering for mesh surfaces*. (2018). doi:10.5281/zenodo.1442584

39. Kimia, B. B. & Siddiqi, K. Geometric Heat Equation and Nonlinear Diffusion of Shapes and Images. *Comput. Vis. Image Underst.* **64**, 305–322 (1996).

40. Lewis, L., Lepage, C., Fournier, M. & Zilles, K. BigBrain: Initial tissue classification and surface extraction. *20th Annual Meeting of the Organization for Human Brain Mapping (OHBM)* (2014).

41. Nair, V. & Hinton, G. E. Rectified linear units improve restricted boltzmann machines. in *Proceedings of the 27th international conference on machine learning (ICML-10)* 807–814

bioRxiv preprint doi: https://doi.org/10.1101/580597; this version posted April 14, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

(2010).

42. Lecun, Y., Bottou, L., Bengio, Y. & Haffner, P. Gradient-based learning applied to document recognition. *Proc. IEEE* **86**, 2278–2324 (1998).
43. Eigen, D. & Fergus, R. Predicting Depth, Surface Normals and Semantic Labels with a Common Multi-Scale Convolutional Architecture. *arXiv [cs.CV]* (2014).
44. Amunts, K. *et al.* BigBrain: an ultrahigh-resolution 3D human brain model. *Science* **340**, 1472–1475 (2013).
45. Fonov, V., Coupe, P., Eskildsen, S. & Collins, D. Atrophy specific MRI brain template for Alzheimer's disease and mild cognitive impairment. *Alzheimers. Dement.* **7**, S717 (2011).
46. Taubin, G. Curve and surface smoothing without shrinkage. in *Proceedings of IEEE International Conference on Computer Vision* 852–857 (IEEE Comput. Soc. Press, 1995).
47. Funck, T. *et al.* Assessing neuronal density in peri-infarct cortex with PET: Effects of cortical topology and partial volume correction. *Hum. Brain Mapp.* **38**, 326–338 (2017).
48. Glasser, M. F. *et al.* A multi-modal parcellation of human cerebral cortex. *Nature* **536**, 171–178 (2016).
49. Elston, G. N., Tweedale, R. & Rosa, M. G. Cellular heterogeneity in cerebral cortex: a study of the morphology of pyramidal neurones in visual areas of the marmoset monkey. *J. Comp. Neurol.* **415**, 33–51 (1999).