
Electronic Thesis and Dissertation Repository

10-23-2020 10:15 AM

Computational Unfolding of the Human Hippocampus

Jordan M. K. DeKraker, *The University of Western Ontario*

Supervisor: Khan, Ali R., *The University of Western Ontario*

Co-Supervisor: Köhler, Stefan, *The University of Western Ontario*

A thesis submitted in partial fulfillment of the requirements for the Doctor of Philosophy degree in Neuroscience

© Jordan M. K. DeKraker 2020

Follow this and additional works at: <https://ir.lib.uwo.ca/etd>



Part of the [Other Neuroscience and Neurobiology Commons](#)

Recommended Citation

DeKraker, Jordan M. K., "Computational Unfolding of the Human Hippocampus" (2020). *Electronic Thesis and Dissertation Repository*. 7443.

<https://ir.lib.uwo.ca/etd/7443>

This Dissertation/Thesis is brought to you for free and open access by Scholarship@Western. It has been accepted for inclusion in Electronic Thesis and Dissertation Repository by an authorized administrator of Scholarship@Western. For more information, please contact wlsadmin@uwo.ca.

Abstract

The hippocampal subfields are defined by their unique cytoarchitectures, which many recent studies have tried to map to human *in-vivo* MRI because of their promise to further our understanding of hippocampal function, or its dysfunction in disease. However, recent anatomical literature has highlighted broad inter-individual variability in hippocampal morphology and subfield locations, much of which can be attributed to different folding configurations within hippocampal (or archicortical) tissue. Inspired in part by analogous surface-based neocortical analysis methods, the current thesis aimed to develop a standardized coordinate framework, or surface-based method, that respects the topology of all hippocampal folding configurations. I developed such a coordinate framework in Chapter 2, which was initialized by detailed manual segmentations of hippocampal grey matter and high myelin laminae which are visible in 7-Tesla MRI and which separate different hippocampal folds. This framework was leveraged to i) computationally unfold the hippocampus which provided implicit topological inter-individual alignment, ii) delineate subfields with high reliability and validity, and iii) extract novel structural features of hippocampal grey matter. In Chapter 3, I applied this coordinate framework to the open source BigBrain 3D histology dataset. With this framework, I computationally extracted morphological and laminar features and showed that they are sufficient to derive hippocampal subfields in a data-driven manner. This underscores the sensitivity of these computational measures and the validity of the applied subfield definitions. Finally, the unfolding coordinate framework developed in Chapter 2 and extended in Chapter 3 requires manual detection of different tissue classes that separate folds in hippocampal grey matter. This is costly in the time and the expertise required. Thus, in Chapter 4, I applied state-of-the-art deep learning methods in the open source Human Connectome Project MRI dataset to automate this process. This allowed for scalable application of the methods described in Chapters 2, 3, and 4 to similar new datasets, with support for extensions to suit data of different modalities or resolutions. Overall, the projects presented here provide multifaceted evidence for the strengths of a surface-based approach to hippocampal analysis as developed in this thesis, and these methods are readily deployable in new neuroimaging work.

Keywords

Hippocampus, Hippocampal subfields, Medial Temporal Lobe, Magnetic Resonance Imaging, 3D Histology, Computational anatomy, Medical image processing, Machine learning

Summary for Lay Audience

One way that scientists have studied the human hippocampus is through Magnetic Resonance Imaging (MRI). While non-invasive, MRI has limitations in resolution and contrast and so the current thesis aims to gain as much information as possible from such images. Given only a blurry image of the hippocampus, we can get a rough idea of its overall shape and properties, which are often estimated in current research and clinical examinations. From existing literature, we know that the hippocampus is composed of a thin, folded sheet of tissue. Thus, there may only be a few possible folding patterns that could produce a given coarse hippocampal shape and image. The current thesis tries to reconstruct possible folding patterns of the hippocampus from MRI or other images.

Once we understand how a given hippocampus is folded, it becomes easier to learn more about its structure. For example, we can measure its thickness and other properties, or we can more easily divide it into contiguous subfields. For easier visualization, we can also computationally ‘unfold’ this structure and plot its properties, such as thickness, across its full 3D extent in just one flattened plane of view. Doing this type of unfolding also allows us to align different hippocampi despite differences in their original, 3D folded shapes. This can be used to align many hippocampi and look for subtle statistical differences between subgroups.

A computational approach to understand hippocampal folding is presented in Chapter 2 of this thesis, but this still relies on some manual input. In Chapter 3 we extend these methods to a unique 3D histological dataset and show that our computational unfolding approach alone can be used to detect hippocampal subfields, instead of more traditional detection by trained neuroanatomists. In Chapter 4 we fully automate the application of this computational unfolding to new MRI data using deep learning instead of manual delineation of tissues separating different folds of the hippocampus. Altogether, these methods could help us

identify which properties of the hippocampus are correlated with performance on tasks, like episodic memory tests, and which properties are correlated with, or diagnostic of, neurological diseases.

Co-Authorship Statement

The work presented in Chapter 2 of this thesis was carried out in collaboration with Kayla Ferko, Dr. Jonathan Lau, Dr. Ali Khan, and Dr. Stefan Köhler, with Ali and Stefan jointly supervising the overall project. The MRI and histology data used in this project was first collected by Ali and colleagues and reused for this project. The coordinate framework was jointly conceived of by Jordan and Ali as a solution to challenges faced by Jordan during manual segmentation. The code was written by Jordan, with input and reused code segments from Ali. Kayla contributed to manual segmentation and to the development of the manual tissue segmentation protocol. Jonathan provided input in interpreting images and contributed to conversations on the development and applications of the hippocampal coordinate framework. Jordan wrote and published a manuscript describing this project in the journal *NeuroImage*, with help from Stefan and revisions by Stefan and Ali.

The work presented in Chapter 3 of this thesis was carried out in collaboration with Kayla Ferko, Dr. Jonathan Lau, Dr. Stefan Köhler, and Dr. Ali Khan, with Stefan and Ali jointly supervising the overall project. Jordan carried out tissue segmentation of the open source BigBrain dataset, while Kayla segmented this tissue into subfields. Jordan extended the coordinate framework code to suit this data, and wrote the feature extraction code. Code for data-driven clustering of these features was written by Jordan, with input from Ali. Jonathan provided input on discussion related to analyses and publication of this work. Jordan wrote and published a manuscript describing this project in the journal *NeuroImage* with revisions by Stefan and Ali.

The work presented in Chapter 4 was carried out under the supervision of Dr. Stefan Köhler and Dr. Ali Khan. Manual segmentations across multiple datasets were provided by Jordan, as well as collected from other projects including the bachelor's thesis of Kayla Ferko, open source published work by Dr. Laura Wisse, Dr. Paul Yuskvevich, *et al.*, and diffusion weighted MRI images segmented by Mohamed Yousif with help from Dr. Roy Haast. Jordan conceived of and wrote the code for U-Net segmentation training, and Sara Pac and Cynthiya Gnanaseelan contributed to evaluation of initial U-Net outputs. Sara and Jordan manually corrected U-Net outputs for use in training on subsequent incremental learning steps. Jordan wrote the code for post-processing of segmentations and subsequent visualization of the

coordinate framework and unfolding. Ali contributed to maintenance of code, including the development of a Brain Imaging Data Standards (BIDS) application wrapper. Jordan will write a manuscript describing this project for publication in the coming months.

Additional contributions to all projects are as follows: Uzair Hussain developed a reparameterization of unfolded hippocampal space according to real-world distances, as well as extensions of this work to diffusion MRI. Roy Haast contributed to discussions of MRI acquisition. Sara Pac contributed to manual segmentation and discussion of the clarity of images and manual protocol steps. Mohamed Yousif helped develop methods to visualize unfolded hippocampal features and contributed to discussions of different MRI methods with respect to the clarity of hippocampal structures. Brad Karat contributed to discussions of extensions of the hippocampal coordinate framework for feature extraction from diffusion MRI. These contributions were formative in the writing of this thesis.

Dr. Ali Khan, Dr. Stefan Köhler, Dr. Roy Haast, and Dr. Jörn Diedrichsen offered significant suggestions and edits to the content and clarity of this thesis.

This work was supported by a Canadian Institutes for Health Research Project Grant (CIHR Grant # 366062) to Dr. Ali Khan and Dr. Stefan Köhler. Jordan was funded through a Natural Sciences and Engineering Research Council doctoral Canadian Graduate Scholarship (NSERC CGS-D).

Acknowledgments

Thanks to Dr. Ali Khan. You've imparted on me a love of coding, maths, and the amazing understanding that they imbue. I especially value the time we spent early on in my graduate career having one-to-one or small group meetings where we worked through problems or brainstormed new combinations of your existing knowledge with new questions. Those meetings opened a whole new domain of knowledge to me. You've radically changed both my direction in research and the tools I use to understand the world around me. Watching your lab grow and branch out has been inspiring to me, and I hope to one day follow in your footsteps.

Thanks to Dr. Stefan Köhler. You've guided me through the turbulent world of academia since my undergrad. Though my topics of interest have jumped around a bit over the years, I think we will always share a broad and deep seeded interest in the phenomenology and the biology of memory. You've helped me grasp concepts and literature that I would never have been able to appreciate on my own. I admire the curious and constructive lab environment you've created, and your diligence which I aim to emulate both socially and in research.

Thanks to Dr. Chris Martin, Dr. Anna Blumenthal, Dr. Jonathan Lau, Dr. Uzair Hussain, Dr. Roy Haast for candid and formative advice, and for being excellent role models.

Thanks to Jason Kai, Haopei Yang, and all members of the Khan and Köhler labs for being such excellent friends and colleagues.

Thanks to all the members and staff of the Brain and Mind Institute and the Robarts Research Institute for creating two supportive and engaging environments that have been a privilege to participate in.

Thanks to my dad, Peter DeKraaker, for his curiosity and enthusiasm for almost any scientific topic and encouraging questions and suggestions for this thesis.

Thanks to my oldest and closest friend, Rowan Pentz, for always challenging me and discussing a lot of science, philosophy, careers, and life together.

Thanks to my fiancée and life partner, Samantha Chin for being there for me from long distance to quarantine, always encouraging me while still keeping me grounded, and patiently listening to, reading, and editing more information about the hippocampus than you ever expected to in a lifetime! You've had to put up with a lot these past few years, and for that I am eternally grateful. You, and our dog Lucy, have made my time in graduate school a truly joyous era in my life.

Table of Contents

Abstract	ii
Summary for Lay Audience	iii
Co-Authorship Statement.....	v
Acknowledgments.....	vii
Table of Contents	ix
List of Tables	xiv
List of Figures	xv
List of Appendices	xxii
Chapter 1	1
1 General Introduction	1
1.1 Motivation.....	1
1.1.1 General interest in the hippocampus.....	1
1.1.2 Issues addressed in this thesis	2
1.2 Hippocampal Anatomy	3
1.2.1 Ontogeny.....	3
1.2.2 Subfields	5
1.3 Hippocampal function.....	7
1.3.1 Subfield-related functions.....	7
1.3.2 Hippocampal long-axis	9
1.3.3 Hippocampal subfields and disease	10
1.4 Neuroimaging of the hippocampus.....	11
1.4.1 Principles of MRI.....	11
1.4.2 MRI of the hippocampus	13
1.4.3 Manual subfield delineation.....	15

1.4.4	Registration-based subfield delineation	19
1.4.5	Classification-based subfield delineation	21
1.4.6	Deep learning-based subfield delineation	22
1.4.7	Surface-based subfield delineation	24
1.5	Goals of this thesis and overview of projects	26
Chapter 2	30
2	Unfolding Coordinate Framework	30
2.1	Introduction.....	30
2.1.1	Goals of the current study	32
2.1.2	Critical structural features we aimed to accommodate	33
2.2	Methods and data	34
2.2.1	Study participants.....	34
2.2.2	MRI acquisition	34
2.2.3	Preprocessing	34
2.2.4	Detection and labelling of the SRLM and hippocampal grey matter	35
2.2.5	Manual subfield segmentation	37
2.2.6	Unfolding of hippocampal grey matter.....	38
2.2.7	Subfield borders in unfolded coordinate space.....	42
2.2.8	Quantitative unfolded tissue properties	43
2.2.9	Histological validation	43
2.3	Results.....	44
2.3.1	Detection and labelling of the SRLM and hippocampal grey matter	44
2.3.2	Subfield borders in unfolded coordinate space.....	44
2.3.3	Quantitative unfolded tissue properties	46
2.3.4	Histological validation	47
2.4	Discussion.....	50

2.4.1	Detection and labelling of the SRLM and hippocampal grey matter	50
2.4.2	Anatomical details of unfolding.....	51
2.4.3	Subfield borders in unfolded coordinate space	52
2.4.4	Quantitative unfolded tissue properties	53
2.4.5	Histological validation	55
2.4.6	Hippocampal unfolding in the context of extant literature	55
2.4.7	Conclusions.....	57
Chapter 3	59
3	Unsupervised 3D Histology Analysis	59
3.1	Introduction.....	59
3.2	Methods.....	62
3.2.1	Materials	62
3.2.2	Manual tracing	63
3.2.3	Topological unfolding framework	66
3.2.4	Morphological feature extraction.....	68
3.2.5	Laminar feature extraction	70
3.2.6	Unsupervised clustering.....	70
3.2.7	Anterior-posterior variation	72
3.3	Results.....	72
3.3.1	Manual tracing	72
3.3.2	Topological unfolding.....	73
3.3.3	Characterization of the hippocampus in unfolded space	74
3.3.4	Unsupervised identification of hippocampal subfields using combination of morphological and laminar features	75
3.3.5	Unsupervised identification of hippocampal subfields using morphological or laminar features in isolation	79
3.3.6	Relative contributions of individual features to subfield clustering	80

3.3.7	Structural variation along the longitudinal hippocampal axis	81
3.4	Discussion	81
3.4.1	Structural characterization of the hippocampus in BigBrain	82
3.4.2	Unsupervised clustering of all features reveals hippocampal subfields ...	83
3.4.3	Morphological features are sufficient to approximate most subfield boundaries	84
3.4.4	Anterior-posterior structural variation	85
3.4.5	MRI applications and future directions.....	87
3.4.6	Data and resources made available	88
3.4.7	Conclusions.....	88
Chapter 4	89
4	Hippocampal Automated Topology	89
4.1	Introduction.....	89
4.2	Methods and data	93
4.2.1	Preprocessing and resampling.....	93
4.2.2	U-Net implementation	93
4.2.3	U-Net training and evaluation.....	94
4.2.4	Post-processing	96
4.2.5	Unfolding and subfield definition.....	98
4.2.6	Comparison to other methods	98
4.3	Results and Discussion	100
4.3.1	U-Net segmentation	100
4.3.2	Post-processing	101
4.3.3	Unfolded subfields and features.....	103
4.3.4	Comparison to other methods	104
4.4	Discussion.....	109

4.4.1	Methodological approach.....	109
4.4.2	Comparison to extant methods.....	111
4.4.3	Limitations and future directions	112
4.4.4	Conclusions.....	114
Chapter 5.....		116
5	General Discussion.....	116
5.1	Background and scope	116
5.2	Summary of projects	117
5.3	Accommodation of specific hippocampal features in the unfolding framework	119
5.3.1	Curvature in the hippocampal head and tail	119
5.3.2	Fine scale folding in the uncus and tail.....	120
5.3.3	Non-topologically defined hippocampal substructures	121
5.4	Adaptation of current approach to new data	123
5.4.1	Higher resolution datasets.....	123
5.4.2	Anisotropic datasets	124
5.4.3	Datasets with T1w or other contrasts.....	124
5.4.4	Diffusion magnetic resonance imaging	125
5.5	Applications	127
5.5.1	Research in patient populations	127
5.5.2	Localization beyond subfields	129
5.5.3	Integration with neocortical surface representations	131
5.5.4	Impact on cognitive neuroscience.....	131
5.6	Conclusions.....	133
References.....		134
Appendices.....		147
Curriculum Vitae		178

List of Tables

Table 1. Volumes of each manually defined subfield (mm ³).	73
Table 2. DSI between k-means clusters and their closest corresponding manually defined subfield.....	79

List of Figures

Figure 1. Topological structure in the hippocampus. A) Illustrates the topological development of the neocortex (beige and purple), archicortex (predominantly hippocampus) (pink, spanned by the green and red arrows) and the remainder of the allocortex as a flattened but continuous cortical sheath. B) Illustrates the inward ‘roll’ of the hippocampus of mammals forming the hippocampal sulcus. C) Illustrates the position and shape of the adult human hippocampus, including digitations and anterior-posterior curvature. In each panel the anterior-posterior (human) or rostral-caudal (rodent or other mammals) dimension of the hippocampus is illustrated by the red arrows, proximal-distal dimension in green arrow, and in panel B) the thickness or inner-outer dimension of the hippocampus is illustrated by the blue arrow. 4

Figure 2. Subfield laminar structure and connectivity. A) Shows a simplified illustration of the laminar distribution of neurons within the different subfields of the hippocampus (DG in red, CA fields in green, subicular complex in blue). B) Illustration of the direct (dashed line) and indirect (solid line) pathways within the hippocampus. The dotted line in CA3 illustrates its high number of recurrent collaterals. 7

Figure 3. Illustration of hippocampal subfield topology in traditional and multiplanar histological slices. A) shows traditional hippocampal slices with and without subfields overlaid. Black lines indicate the continuity of subfields from the dentate gyrus to medial temporal lobe neocortex. B) shows the approximate locations of the sagittal slice shown (outlined in orange) in the full brain, and coronal slices (outlined in solid or dashed blue lines) on the left 3D model of the hippocampus. In particular, slices through the hippocampal head (light blue dashed line) and tail (dark blue dashed line) show within-plane discontinuities. C) shows multi-planar resampled slices along the axis of hippocampal curvature with and without subfields overlaid. The colour of each image outline corresponds to the location of its sampling on the right 3D hippocampal model in B). The continuity of subfields is again shown by a black line and is consistent in all slices in C). 17

Figure 4. Variability in overall hippocampal shape in four representative individuals as measured in MRI. 3D models are shown to the left of each grey box, with a coronal slice

through the hippocampal head (outlined in blue) and a sagittal slice (outlined in orange) to the right. Hippocampi were affine registered such that all slices shown here are from equivalent positions. Hippocampi were organized along two arbitrary, qualitative axes: curvature and digitation. Curvature refers to the long-axis, which can appear more ‘J’ shaped with a straight hippocampal body, or more ‘C’ shaped with a curved hippocampal body. This is most easily seen in the 3D model. Digitation refers to the number and amplitude of grey matter folds within the hippocampus, and can be seen on the 3D model, the coronal plane when digitations are in the hippocampal head, and the sagittal plane when digitations are in the hippocampal body. Equivalent coronal planes can show dramatically different subfield organization depending on these morphological differences..... 18

Figure 5. Equivalent coronal slices showing the left collateral sulcus in two Human Connectome Project subjects (Elam and Van Essen, 2015) (minimally preprocessed and parcellated using Freesurfer’s surface-based APARC neocortical parcellation). Labels include hippocampus (orange; ASEG subcortical parcellation method), parahippocampal cortex (green), fusiform cortex (sage) and inferior temporal cortex (pink)..... 25

Figure 6. Illustration of SRLM (green) and hippocampal grey matter (red) labelling. A) 3D model of the SRLM label in the center, with representative coronal slices from the head and body on the sides, at the positions indicated by the dotted lines. B) Same views as above but depicting the SRLM label after spherical dilation. C) Same views as above showing the combined SRLM label and grey matter label, after manual adjustments to the grey matter label..... 35

Figure 7. Representative slices showing manual subfield labelling from hippocampal head (top) and body (tail). Left shows reference materials from (Ding and Van Hoesen, 2015), middle shows manually traces SRLM (green), hippocampal grey matter (red) and cysts (yellow), right shows manually delimited subiculum (brown), CA1(blue), CA2 (white), CA3 (pink), and dentate gyrus (cyan). See Appendix 2 for further details. Images were adapted with permission from (Ding and Van Hoesen, 2015). 38

Figure 8. Illustration of Laplacian unfolding along the long-axis, proximal-distal, and laminar dimensions in A), B), and C), respectively. The upper left inset image in A) shows a 3D model of the SRLM (green) and grey matter (red) labels, with the HATA (pink) and

indusium griseum (grey) to be used as boundaries for Laplace's equation. The lower left image in A) shows arbitrarily coloured bins within the resulting potential field gradient. To the right is the same model as the lower left but showing cross sections from the body (top) and head (lower three), depicting in particular the uncus (lower two) and vertical component of the uncus (bottom). The locations of these cross sections are shown by the black dotted lines (1–3) and the eye figure (4). B) shows the same views of the same hippocampus but using the inner dentate gyrus (turquoise) and medial temporal lobe cortex border (white) (upper left insert in B), as boundaries for Laplace's equation for the proximal-distal dimension. C) shows the same views of the same hippocampus but using the SRLM (seen in green under semi-transparent red grey matter) and outer hippocampal borders (upper left insert in C) as boundaries for Laplace's equation for the laminar dimension. White dotted lines in C) (right) show the true laminar structure of CA4 and DG, which is not respected in our laminar potential field gradient..... 40

Figure 9. Hippocampal segmentations in unfolded coordinate space. A) Example of a manual subfield segmentation based on (Ding and Van Hoesen, 2015). Similarly, B) shows an exemplar with fewer digitations and less medial curvature. C) shows ‘unfolded’ hippocampal grey matter, with subfield label identity determined at each long-axis and proximal-distal coordinate from the manual segmentations (winner-takes-all over the sample). The shaded areas indicate standard error of the mean for each subfield boundary location across the sample of manual segmentations. Dotted lines approximately indicate commonly used boundaries between the hippocampal head, body, and tail, with the head further subdivided into uncus and vertical component of the uncus. 45

Figure 10. Spatial overlap in Dice Similarity Index (DSI) between manual subfield segmentations (manual-manual) and between leave-one-out unfolded group-average subfield segmentations and their manually segmented counterparts (unfolded-manual). The leave-one-out technique was performed such that borders from one participant's left and right hippocampi were not included in the averaged borders that informed unfolded segmentation of that participant's hippocampi. 46

Figure 11. Quantitative mapping in unfolded coordinate space compared to subfield borders. A) Average intracortical myelin estimates (T1 over T2 MR intensities; arbitrary units). B)

Average cortical thickness. Both A) and B) have average subfield borders overlaid. Note that in the dentate gyrus, thickness estimates are perpendicular to the true laminar structure (see Chapter 2.3.2). C) Manual and unfolded subfield segmentations compared to intracortical myelin in highly digitated (top) and less digitated (bottom) representative hippocampal models. Average intracortical myelin is mapped to the surface models of these hippocampi for easier comparison. 47

Figure 12. Neu-N stain of resected hippocampal tissue with comparison of histologically segmented hippocampal subfields to *in-vivo* subfield labels and unfolding Laplace gradients in the same individual. Slices move from posterior (hippocampal body) to anterior through the hippocampal head and are separated by 4.4mm. Red arrows indicate marked cell loss in distal CA1. The proximal-distal and long-axis gradients are surrounded by low colour-scaled voxels because of their interpolation when transforming to the histological space. 48

Figure 13. Manual traces of hippocampal archicortex and segmentation into subfields. A) shows coronal slices through the left hippocampal head (rows 1–3) body (row 4) and tail (row 5), with manual segmentations overlaid in the images to the right. Images were taken from coronal slices 716, 632, 590, 376, and 230 of the 40 µm native space left hippocampal block. B) shows 3D models of each hippocampus as seen from their superior aspect, with the inferior aspect shown in C). Dotted lines in B) indicate approximate locations of each coronal slice shown in A). SRLM, vestigial hippocampal sulcus, alveus, and fimbria were excluded from all labels. Red arrows indicate anterior folding in the vertical component of the uncus, the orange inset in the third row highlights ‘islands’ of neuronal cell bodies in the subicular stratum lacunosum, and yellow arrows indicate gyrifications in the posterior body and tail of the hippocampus. 64

Figure 14. Topological unfolding framework in BigBrain with hippocampal subfields. A) Sagittal slice and 3D models of the Laplacian solutions (proximal-distal and anterior-posterior) for the right hippocampus. Image was taken from sagittal slice 514 of the 40 µm native space left hippocampal block. B) Mid-surface topological models of the left and right hippocampi in native and unfolded space. 67

Figure 15. Characterization of the hippocampus using morphological and laminar features. The top diagrams illustrate how each feature is derived (see Chapters 3.2.4 and 3.2.5 for

details). Top left shows an example segment of cortex, while the top center and top right show an example laminar profile and its absolute derivative (Abs.Deriv), respectively (Amunts *et al.*, 1999; Larsen, Nielsen and Sporring, 2006). Heat maps below show the z-scored values of each feature across the unfolded hippocampus in the left and right hemispheres, with the same colour scaling in both hemispheres. Overlaid in white are manually defined subfield borders, the top edge being the border with the DG which is out-of-plane. 69

Figure 16. Unsupervised k-means clustering of features. The left images show k-means clusters in unfolded space at $k = 5$, with manually defined subfield borders overlaid in white. The right images show the same data in native space, with 10% anterior and posterior edges extrapolated by nearest neighbour. Clustering was completed for the combined set of all features, laminar features only, and morphological features only. 77

Figure 17. Exploration of inherent feature variance. A) shows PCA component loadings from each feature with a dotted line at 1% after which subsequent components were discarded. B) shows the feature loadings of the first 8 components, with multiple rows for the various smoothing kernels applied to each feature. C) shows the correlation between all features, with separate boxes around morphological and laminar features. D) shows a visualization of the first 8 components, with manually defined subfield borders overlaid in white. 78

Figure 18. Features of interest plotted with respect to the anterior-posterior axis of the hippocampus. Colours indicate manually defined subfields, and shaded areas indicate standard deviation. Data are combined across the left and right hippocampi. a.u. stands for arbitrary units, see Chapter 3.2.7 for additional details. 78

Figure 19. Overview of full Hippocampal Automated Topology pipeline. Step 1: left and right hippocampi are cropped and resampled to 0.3mm^3 isotropic obliquely to the hippocampal long-axis. Step 2: images are segmented via U-Net architecture and then post-processed using template shape injection. Step 3: Previously developed Laplace coordinate framework is applied to the domain of hippocampal grey matter tissue. Step 4: A single subfield atlas defined in 2D unfolded space, in this case generated from BigBrain 3D histology, is propagated to a given subject's native hippocampal folding configuration. Shown here is the example of the right hippocampus of HCP1200 subject 108020. 92

Figure 20. Benchmarking of tissue segmentation following the final incremental U-Net training step. A) shows the training of the final U-Net incremental learning step over 500 000 iterations. Solid colours are smoothed over iterations and faded colours indicate raw loss at each iteration. B) shows the Dice overlap between manual and automatically generated segmentations for each label of interest in the test set (15%). Results are shown both before and after post-processing via template shape injection..... 100

Figure 21. Effect of template choice and update field variance (a parameter regulating fluidity) on template shape injection results. Only GM (red) and SRLM (orange) are shown for clarity. The original tissue segmentation contains topological breaks (green arrows) due to segmentation errors, which are problematic for subsequent unfolding. Highly fluid template shape injection from an averaged hippocampal shape, as well as several notably different exemplars are shown. The effect of template choice is particularly notable at lower update field variance. Final parameter choices are indicated by the black box. Segmentation errors persist in this case (purple arrows), but do not include topological breaks..... 102

Figure 22. Average features across all HCP1200 subjects plotted in unfolded space. A group-averaged mid-surface in native space is shown in the top left..... 104

Figure 23. Bland-Altman plots comparing total hippocampal volume between FS, HAT, and ASHS. Units are in mm^3 , and volumes include both hippocampal grey matter and SRLM. 105

Figure 24. Subfield definitions using the method proposed here Hippocampal AutoTop (HAT), Freesurfer6.0 (FS), and Automatic Segmentation of Hippocampal Subfields (ASHS). The first column shows a single subject example (HCP1200, 105923) subfield definitions using each method. The second column shows tables of the average subfield volumes using each method. The third column shows the Dice overlap between equivalent subfields from each method. The final column illustrates the equivalency between methods. 106

Figure 25. Deformations between equivalent subfield produced using the method proposed here Hippocampal AutoTop (HAT), Freesurfer6.0 (FS), and Automatic Segmentation of Hippocampal Subfields (ASHS). A) shows equivalent labels from each method from a single subject, as well as the fluid label-label 3D deformations between them. Shown here is the example of the left hippocampus of HCP1200 subject 105923. B) shows the group-averaged

deformations from a HAT to each other method (right), as well as the mode subfield identities projected into unfolded space and on a group-averaged mid-surface (left)..... 108

Figure 26. Mid-surface models of hippocampal folding estimated through MRI (A) and 3D histology (B). Purple arrows indicate folds seen in histology that are not typically seen using *in-vivo* MRI, while black arrows represent gross curvature and digitations that are typically seen in both modalities. Digitations in the hippocampal body are not always present and are absent from this MRI sample. Such digitations are typically visible but likely attenuated in MRI under current acquisition protocols. 119

Figure 27. Sample scans employing T2w, T1w, and low b-value dMRI acquisition protocols. The T2w scan represents a different subject, since no directly comparable data were available. Note that CSF and white matter within and surrounding the hippocampus show opposite intensities in T2w and T1w images, but not in dMRI images..... 125

Figure 28. Extension of unfolding (or surface-based) coordinate framework to surrounding medial temporal lobe neocortical structures. A) shows an extension of our proximal-distal hippocampal axis in the hippocampus (yellow) into neighbouring neocortex (orange, red, then black) according to geodesic distance. B) shows an example of a statistical analysis (based on singular value decomposition) that employs contiguous points rather than distinct parcels or ROIs. Data from (Paquola *et al.*, in submission). 130

List of Appendices

Appendix 1 (Chapter 2): Manual and user-guided steps for labelling of SRLM, hippocampal grey matter, and extrahippocampal structures for subsequent unfolding	147
Appendix 2 (Chapter 2): Manual subfield segmentation by alignment with Ding & Van Hoesen, 2015; Duvernoy <i>et al.</i> , 2013	161
Appendix 3 (Chapter 2): Manual histology labelling protocol	164
Appendix 4 (Chapter 3): Additional anatomical details of the hippocampus and surrounding structures in BigBrain	167
Appendix 5 (Chapter 3): Additional methodological details	169
Appendix 6 (Chapter 3): Laminar modelling using Laplacian vs. Equivolume model in BigBrain	171
Appendix 7 (Chapter 3): Additional clustering results	172
Appendix 8 (Chapter 4): Acquisition and preprocessing details of datasets used in U-Net training	176

Chapter 1

1 General Introduction

1.1 Motivation

1.1.1 General interest in the hippocampus

On an evolutionary scale, the archicortex (specifically, the subcomponent now known as the hippocampus) lies at an information processing junction. Prior to its phylogenetic development, the nervous system followed a relatively rigid mapping of stimulus sets to response patterns. This was predominantly encoded in genetics and manifested through the ontological development of the nervous system. After the advent of the hippocampus (or its analogues in non-mammalian species) a new era of behavioural flexibility began to emerge. That is, many behaviours were shaped to a greater extent by the environment rather than genetically coded. Some of these behaviours and requisite cognitive operations that have been discussed in recent literature include context-dependent behavioural repertoires (Pravosudov and Roth, 2013), navigation (O'Keefe and Nadel, 1978), episodic memory (Tulving and Markowitsch, 1998), transitive inference (Dusek and Eichenbaum, 1997), episodic future-thinking (Schacter, Benoit and Szpunar, 2017), and reorganization of sensory, semantic, or abstract concept representations along continuous 'pseudo-spatial' dimensions (Garvert, Dolan and Behrens, 2017). The hippocampus is also the evolutionary precursor to the more cytoarchitecturally complex neocortex, a structure whose expansion coincides with increased abilities in nearly all cognitive domains. Concurrently with expansion in the neocortex, hippocampal functions appear to be maintained, including its role in the acquisition of complex new behaviours and memories (Squire, 2009; Eichenbaum, 2013). Thus the operations carried out in the hippocampus are characterized by some unique properties not seen anywhere else in the brain; in particular these operations seem to include fast and flexible acquisition of spatial or non-spatial representations, which may later become consolidated throughout the neocortex (McClelland, McNaughton and O'Reilly, 1995).

The hippocampus is not only critical for a wide variety of functions, but it is also uniquely sensitive to a variety of diseases and, more generally, metabolic insults. Worth highlighting are hypoxia, focal epilepsy, Alzheimer's disease, clinical depression, and schizophrenia, though many other conditions have also been shown to be associated with lack of hippocampal integrity more so than the integrity of the brain as a whole. Some broad factors that may contribute to hippocampus-specific tissue abnormalities in these disorders include: i) the high metabolism of the hippocampus relative to other brain structures, making metabolic insults potentially more damaging (Small *et al.*, 2011), ii) the predominance of large, highly plastic pyramidal neurons from which new adaptive or maladaptive microcircuits can quickly emerge (Malenka and Bear, 2004), iii) the fact that the hippocampus is the primary site of human adult neurogenesis (Kozareva, Cryan and Nolan, 2019), and iv) the many neuromodulators affecting the hippocampus, including high levels of acetylcholine (Menschik and Finkel, 1998), serotonin (Gould, 1999), noradrenaline (Madison and Nicoll, 1982), and corticosterone (Diamond *et al.*, 1992) compared to other brain structures. This provides many possible ways that hippocampal activity could become dysregulated when these neuromodulators are dysregulated. Many researchers have also proposed links between specific diseases and hippocampal subfields and functions, which will be reviewed in Chapter 1.3.

1.1.2 Issues addressed in this thesis

With these motivations in mind, the current thesis will explore a computational framework for studying the 3D anatomy of the hippocampus, which can be used in future work to improve our understanding of its relationship to function and disease. Briefly, this framework entails modelling hippocampal tissue as a folded cortical surface, which can improve our understanding of this structure over traditional methods that consider only its 3D volume or 2D area as a whole. The methods used will span 3D *ex-vivo* histology at a microscale (i.e. on the scale of cortical laminae and cell assemblies) to *in-vivo* and *ex-vivo* neuroimaging using Magnetic Resonance Imaging (MRI) and the mesoscale (i.e. on the scale of gyri, sulci, and cortical columns including hundreds of neurons). The development of this coordinate framework, initially in 7-Tesla MRI data, is presented in Chapter 2. In Chapter 3 we apply this framework to 3D histology with

microscopic detail and show that features extracted using this framework are sufficient to derive the known hippocampal subfields. In Chapter 4 we fully automate the application of this framework and subfield segmentation to *in-vivo* MRI data using deep learning methods.

As background to this exploration, I will overview hippocampal anatomy, ontogeny, and its division into subfields. To further motivate this line of investigation, and its applications, I will examine several examples of prominent theories on the links between hippocampal subfields, function, and implication in disease. Finally, I will overview current methods for imaging and analyzing the hippocampus, including acquisition methods in MRI and extant segmentation methods.

1.2 Hippocampal Anatomy

1.2.1 Ontogeny

The brain lies at the topological anterior of the embryonic neural tube - a cylindrically organized set of neuronal precursors. Extensions from this neural tube innervate the different body segments, including major organs and appendages where neural pathways terminate in sensory or motor afferents or efferents (Sperber, 1995). A similar topologically ordered organizing principle holds true in the brain, where a huge expansion in the number of neurons is seen, causing the neural tube to ‘bubble’ outwards. This ‘bubbling’ provides 2D topological organization to what might otherwise appear to be a chaotic arrangement of tissue and neurons. Though the details of this organization are still debated (see Puelles *et al.*, 2019), the archicortex and specifically the dentate gyrus subfield of the hippocampus mark a segment of the true terminus to the topologically continuous cortex of the brain, as illustrated in Figure 1A.




Figure 1. Topological structure in the hippocampus. A) Illustrates the topological development of the neocortex (beige and purple), archicortex (predominantly hippocampus) (pink, spanned by the green and red arrows) and the remainder of the allocortex as a flattened but continuous cortical sheath. B) Illustrates the inward ‘roll’ of the hippocampus of mammals forming the hippocampal sulcus. C) Illustrates the position and shape of the adult human hippocampus, including digitations and anterior-posterior curvature. In each panel the anterior-posterior (human) or rostral-caudal (rodent or other mammals) dimension of the hippocampus is illustrated by the red arrows, proximal-distal dimension in green arrow, and in panel B) the thickness or inner-outer dimension of the hippocampus is illustrated by the blue arrow.

In mammals, the hippocampus rolls medially upon itself during development, creating the hippocampal sulcus (Nieuwenhuys, Voogd and van Huijzen, 1981; Duvernoy, Cattin and Risold, 2013). Simultaneously, the distal-most edge detaches from the rest of the archicortex and forms a distinct structure known as the dentate gyrus (DG), which wraps itself around the edge of the hippocampus forming two interlocking ‘U’ shapes (see Figure 1B) along its long-axis. In humans, the entire hippocampus migrates into the medial temporal lobes, rotating it from the dorsal-ventral positioning observed in rodents to an anterior-posterior orientation in humans. The trajectory of this migration can be traced along the indusium griseum, which extends from the hippocampus along the cingulum, sitting below the cingulate cortex which is where the hippocampus is found

in rodents. In different mammalian species the hippocampus can be found somewhere in between this track, depending on how distinct that species' temporal lobe is from the rest of the brain.

In humans, the morphology of the hippocampus is further complicated by folding along its anterior-posterior axis (Figure 1C). This complex folding is most pronounced in the anterior portion, which curves medially, then posteriorly, and then superiorly (Ding and Van Hoesen, 2015). This section is often referred to as the uncus of the hippocampus, and the component of the uncus that curls superiorly is known as the vertical component. I will generally refer to this gross folding as the long-axis curvature of the hippocampus. Further folding occurs along the rest of the hippocampus, which is sometimes referred to as digitation, dentation, or gyrification (I will use the term digitation) (Ding and Van Hoesen, 2015; Chang *et al.*, 2018). The hippocampus is also often further divided into head, body, and tail along its longitudinal axis, but these terms do not mark clearly defined structural distinctions (though functional and connectivity differences are also described in section 1.3.2). Note that none of these terms refer to cytoarchitectonically distinct tissue. Instead, they are used colloquially to divide the hippocampus from anterior to posterior, and borders are typically defined heuristically (as in the dashed lines in Figure 1C).

1.2.2 Subfields

The archicortical tissue of the hippocampus is further differentiated during its development into cytoarchitectonically distinct subfields (Nieuwenhuys, Voogd and van Huijzen, 1981; Duvernoy, Cattin and Risold, 2013). The DG lies at the distal edge, followed by cornu ammonis (CA) fields 4 to 1 and the subicular complex (Sub). Sub is sometimes further divided into pre/prosubiculum, subiculum proper, perisubiculum, and parasubiculum (Ding, 2013; Ding and Van Hoesen, 2015). Parasubiculum is continuous with the neocortex, specifically the entorhinal cortex and, more posteriorly, parahippocampal cortex. These subdivisions are what researchers most commonly mean by the 'hippocampal subfields', but there is not currently universal agreement on the labels or precise boundaries for these subregions in either the neuroimaging or in the histology literatures (Yushkevich, Amaral, *et al.*, 2015). There is an ongoing effort to

generate an internationally harmonized set of labels and boundary definitions to improve reproducibility and reduce confusion in the neuroimaging literature (Wisse *et al.*, 2017; Olsen *et al.*, 2019).

The definition of these subfields was originally based on histological features and, particularly, the laminar organization of neurons. In brief, the neocortex contains six unique laminae, typically referred to as layers I to VI. The hippocampus is simpler in structure, and ranges between three laminae in the DG to something closely resembling the full six neocortical laminae in the most proximal regions of the subiculum (e.g. parasubiculum) (see Figure 2A and Ding (2013) for more details). It should also be noted that white matter structures like the alveus, or stratum radiatum, lacunosum, and moleculare (collectively SRLM, also sometimes referred to as the ‘dark band’ in MRI) are often considered to be laminae of the hippocampus despite primarily containing axons and dendrites and few or no neurons. Thus, the number of laminae within the hippocampus can vary by subfield and by the nomenclature that is applied.

1.2.3 Subfield connectivity

Each subfield also has a unique connectivity profile, which is summarized in Figure 2B. Briefly, hippocampal connectivity with the neocortex consists of a highly recurrent circuit with primary inputs and primary outputs in the entorhinal cortex (specifically in the superficial and deep entorhinal cortical layers, respectively). This can be seen in the indirect (also known as the trisynaptic) pathway, and additionally in the direct (or monosynaptic) pathway which follows a similar trajectory but bypasses some subfields. Hippocampus and entorhinal cortex are often considered the apex of the ventral visual stream, a view that is sometimes referred to as the representational-hierarchical perspective (Murray, Bussey and Saksida, 2007). This is in part because of connections from the entorhinal cortex to lower order visual areas like perirhinal, parahippocampal, and lateral occipital cortex. However, additional connections also exist between the entorhinal cortex and the amygdala, prefrontal cortex, and nucleus accumbens. To some extent, the regions of the entorhinal cortex projecting to different brain structures may maintain differential connectivity with CA1 and subiculum subfields, the output regions

of the hippocampus. Thus, these sets of hippocampal and entorhinal connections may even make up distinct or partially overlapping functional subsystems, as described by Aggleton (2012). Not shown in this Figure 2 is a secondary output that projects to the anterior thalamic nuclei, mammillary bodies, and retrosplenial cortex via the white matter structures alveus, fimbria, fornix, and hippocampal commissure. Several other additional connections also exist, including afferents from numerous neuromodulators (see Duvernoy, Cattin and Risold, 2013).

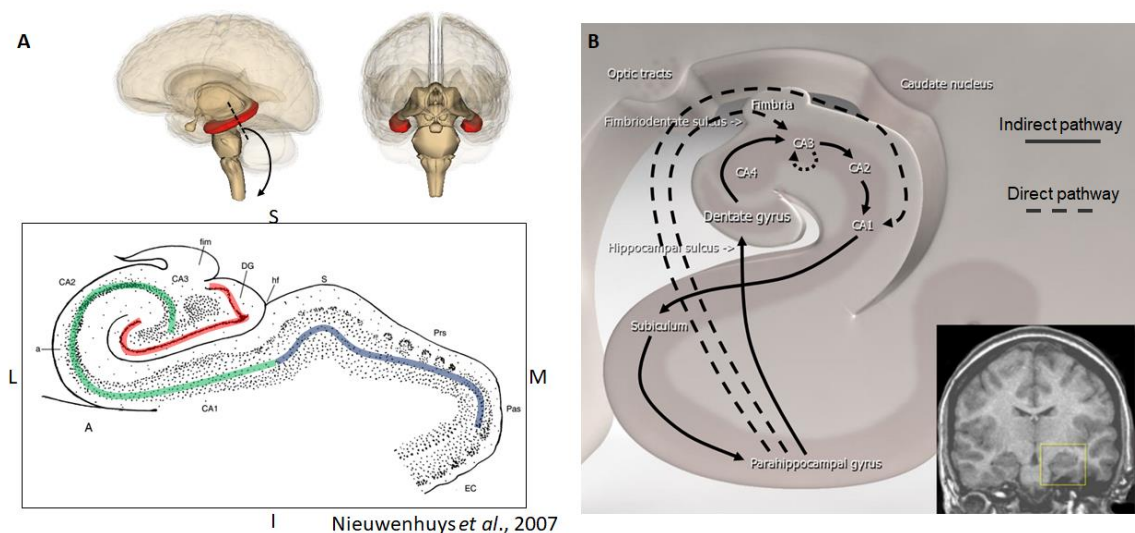


Figure 2. Subfield laminar structure and connectivity. A) Shows a simplified illustration of the laminar distribution of neurons within the different subfields of the hippocampus (DG in red, CA fields in green, subicular complex in blue). B) Illustration of the direct (dashed line) and indirect (solid line) pathways within the hippocampus. The dotted line in CA3 illustrates its high number of recurrent collaterals.

1.3 Hippocampal function

1.3.1 Subfield-related functions

Subfields represent a research avenue to link stereotyped cytoarchitecture to functions that can be described computationally. That is, this line of research aims to mechanistically align intrahippocampal circuitry characteristic of each subfield to

functions observable through behavior and cognition, but it should still be noted that other theories have also garnered support based on observations or manipulations of the hippocampus as a whole rather than individual subfields. Many recent theories of hippocampal function build on the unique cytoarchitecture and connectivity of each of the hippocampal subfields, and this thesis will overview only a few examples. One example is an emerging literature on ‘pattern separation’ and ‘pattern completion’ in the hippocampus (Marr, 1971; Hunsaker and Kesner, 2013; Rolls, 2016). Pattern separation can be defined mathematically as the orthogonalization of overlapping inputs into distinct outputs. This can be achieved by having a relatively small number of neurons project to a relatively large number of neurons with few overlapping connections (also called expansion recoding) (Marr, 1971). In other words, two representations may share some overlapping neurons in a small region A, but they are unlikely to share overlapping neurons after projecting to larger region B. The DG is uniquely well suited to perform this operation, since it consists of many densely packed neurons and receives inputs from a relatively small population of neurons in the entorhinal cortex. The number of neurons in these two regions differs by several orders of magnitude (despite the fact that the DG is volumetrically smaller than the entorhinal cortex), and there are few lateral connections between neurons within the DG. This characterization also fits well with the dentate gyrus’ position in the indirect hippocampal pathway: pattern separated information about an ongoing event may be passed from the dentate gyrus to the CA fields, and modification to the synapses on the CA fields and from the CA fields onward might ultimately support the highly distinctive nature of episodic memory (Aimone, Deng and Gage, 2011; Hunsaker and Kesner, 2013; Kyle *et al.*, 2015; Liu *et al.*, 2016; Rolls, 2016; Rolls and Kesner, 2016; Leal and Yassa, 2018).

The complementary function to pattern separation is pattern completion, the retrieval of a stored representation given only a partial or degraded cue. This type of computation has been extensively studied *in-silico* via computational modelling, and is often referred to as attractor dynamics within an auto-associator network (Knierim and Zhang, 2012), which has also received recent interest in machine learning (e.g. Bakiras *et al.*, 2015; Chappell and Humphreys, 1994). Conceptually, the attractor substrate, or network, will recurrently modify its activity until it reaches some stable state. The stable state will depend on the

inputs from which the attractor network was initialized, as well as its previously structured internal connections. With small to moderate changes in the input signal, the attractor will still produce the same output (i.e. completing the pattern that was present at encoding). However, beyond some threshold of changes to the input, the auto associative attractor will begin to form other states and thus produce a very different output, which could correspond to another previously encoded stable state or a completely novel one. This operation has been ascribed to CA3, which has a high number of recurrent collaterals (Figure 2B) making it well suited to act as an auto associative attractor network.

Pattern separation and pattern completion have received much investigation in recent studies (Liu *et al.*, 2016), and may be involved in a wide array of functions beyond their roles in episodic memory. For example, hippocampal pattern separation and completion have been implicated in real-time navigation and perceptual discrimination. In navigation, selection of the correct memory of the layout of an environment among competing, highly similar memories may rely on both hippocampal pattern separation and completion (Kyle *et al.*, 2015). In perception, hippocampal pattern completion may help guide predictions about a visual stimulus (Hindy, Ng and Turk-Browne, 2016) while pattern separation may facilitate discrimination of highly similar stimuli (Lee, Yeung and Barense, 2012). Pattern separation and completion represent one computational framework for understanding some functions of hippocampal subfields DG and CA3, but many other theories exist to describe functions of hippocampal subfields or of the hippocampus in its entirety (see Moscovitch *et al.*, 2016).

1.3.2 Hippocampal long-axis

In addition to differences in hippocampal function based on subfield cytoarchitecture, differences in representational content have been proposed along the anterior-posterior (i.e., head-tail) axis of the hippocampus. There are many competing and complementary theories of anterior-posterior specialization of the hippocampus, so this thesis will simply summarize some of the predominant literature. The anterior hippocampus, antero-lateral entorhinal cortex, and perirhinal cortex are closely connected, and are associated with object-based holistic representations at the apex of the ventral visual stream. Conversely,

the posterior hippocampus, postero-medial entorhinal cortex, and parahippocampal cortex are also closely connected and are associated with spatial and perceptually detailed representations (Colombo *et al.*, 1998; Poppenk *et al.*, 2013; Strange *et al.*, 2014; Zeidman and Maguire, 2016; Plachti *et al.*, 2019). Some of these theories further suggest that anterior hippocampus processes primarily transmodal or abstract representations, which may be linked to its stronger connectivity with inferior frontal cortex, whereas posterior hippocampus processes more visuo-spatial representations (Save and Poucet, 2000; Possin, 2010). The bulk of the evidence for these theories comes from *in-vivo* and *ex-vivo* connectivity studies as well as neuroimaging studies aiming to passively quantify representational contents or manipulate tasks demands in one direction or the other while observing hippocampal activity. Congruently, recent work still under review has found continuously graded differences along the long-axis of the hippocampus in measures of connectivity with other widespread brain regions as well as stereotyped brain networks characteristic of rest (or default mode), visuospatial processing, and other high-order processes (Paquola *et al.*, in submission). Through these studies, we get an overall picture of the apparent representational contents of the anterior and posterior hippocampus. However, it is important to note that many intermediate regions contain representational information from both and anterior and posterior connectivity networks, and it may even be that the mixing of these signals in some way contributes to the functions of the hippocampus and perhaps the richness of episodic memories it supports.

1.3.3 Hippocampal subfields and disease

Many neurological diseases and disorders affect the hippocampus, or even specific hippocampal subfields, preferentially over other brain structures. One salient example comes from medial-temporal lobe epilepsy (mTLE), where *ex-vivo* examination of hippocampal subfields has given rise to mTLE subtypes that are predictive of treatment outcomes (Blümcke *et al.*, 2013; Blumcke *et al.*, 2017). Specifically, Type 1 consists of neuronal loss primarily in CA1 and CA4, Type 2 consists of primarily CA1 neuronal loss, and Type 3 consists of primarily CA4 neuronal loss. Type 1 appears to show particularly good prognosis following resective surgery, and thus determining patients' mTLE subtypes can have significant value for planning treatment. Though typically carried out

using *ex-vivo* or resected tissue, this subtyping framework is beginning to show clinical value using *in-vivo* neuroimaging (Steve *et al.*, 2020).

Clinical depression, or major depressive disorder (MDD), has also been tightly linked to decreases in whole hippocampal volume as well as decreased DG and CA subfield volumes (Maruszak and Thuret, 2014; Malykhin and Coupland, 2015). Several hypotheses exist to explain this link mechanistically involving glucocorticoid receptors which are particularly dense in the hippocampus (Na *et al.*, 2014), neuroplasticity (Bianchi, Hagan and Heidbreder, 2005), and some studies have even linked MDD to reduced adult neurogenesis within the DG (Sahay and Hen, 2007). Recovery from depression following treatment with common antidepressant drugs (selective serotonin reuptake inhibitors) or, in intractable cases, following electroconvulsive therapy (ECT) has also been linked to increased hippocampal volumes (Nordanskog *et al.*, 2010; Boldrini *et al.*, 2013). Another context in which hippocampal volume increases have been reliably observed is following aerobic exercise over time periods as short as weeks (Firth *et al.*, 2018). This topic has received much scientific interest and generated excitement in the media in recent years for the prospect of growing new neurons into adulthood which was previously thought to be impossible, but it should be noted that there are many possible causes of hippocampal volume increases, including synaptogenesis, neurogenesis, angiogenesis, or general inflammation. Nevertheless, this line of research is highly promising, especially using *in-vivo* neuroimaging methods or even behavioural testing targeted at hippocampal subfields (e.g. Clelland *et al.*, 2009; Liu *et al.*, 2016; Leal and Yassa, 2018).

1.4 Neuroimaging of the hippocampus

1.4.1 Principles of MRI

MRI is the dominant method for collecting *in-vivo* structural brain data since it is minimally invasive, has reasonable spatial resolution (on the order of millimeters), and can achieve different types of contrast to visualize different tissue properties. However, MRI does have limitations in the quality of images that can be collected. This can be explained by some of the basic principles of MRI. Protons with magnetic moments

(mostly hydrogen) align their spins to a constant primary magnetic field (B_0), and an electromagnetic radio-frequency pulse (RF pulse) is applied, which disrupts proton alignment to B_0 . As protons relaxate to their initial alignment to B_0 , they emit their own echo RF pulses. This realignment can additionally be measured as decay in magnetization in the axis transverse to B_0 (T2-weighted imaging), or, conversely, return of magnetization longitudinal to B_0 (T1-weighted imaging). The period between consecutive RF pulses (or repetition time, TR) and the time between a RF pulse and the response signal measurement (or echo time, TE) can be parameterized in order to highlight differences between common tissues such as fat (high in T1-weighted imaging) or water (high in T2-weighted imaging). Cerebrospinal fluid is primarily water, white matter is composed of heavily myelinated axons with high fat content, and grey matter consists of a mix (where cell bodies and interstitial fluids tend to have more water while axonal or dendritic structures tend to contain more fat in the form of cell membranes). Other related MRI acquisitions leverage features such as the distortion of B_0 caused by deoxygenated iron in blood (the basic principle of T2* or functional MRI), or the movement of water molecules between the time of their perturbation and return to spin alignment with B_0 (the basic principle of diffusion weighted MRI or dMRI) (Mitchell and Mark Cohen, 2004).

MRI signals can be localized in 3D by introducing additive gradients to B_0 in the x, y, or z directions, which is overviewed by Hashemi *et al.* (2010). Briefly, protons have a given precession frequency following the Larmor equation: $f = \gamma * B$, meaning that they will only optimally interact with RF pulses at a certain frequency that varies with magnetic field strength, B. Thus, introducing a gradient in B allows us to select slices that can be optimally excited by RF pulses with the corresponding frequency. This is typically referred to as the z or slice direction. By briefly introducing a B gradient pulse in the y-direction after an RF pulse, we can then cause a phase shift in the subsequent precession of excited protons. Finally, at the time of the RF echo we can introduce a gradient pulse in the x-direction which will modify the frequency of precession. Altogether, the carefully timed introduction of gradients in the z, y, and x directions can be used to relate proton relaxation to a given slice, phase, and frequency. Phase and frequency are referred to as K-space and can readily related to x and y via the Fourier transform, outputting data

that is binned into discrete 2D pixels which can then be stacked in 3D. Alternatively, with more sophisticated gradient pulse designs, data can be acquired in 3D K-space and Fourier transformed to a 3D image, but I will not cover such methods in this thesis.

There is a trade-off between the volume of a given voxel (or resolution) and the amount of signal that can be reliably detected against noise from other sources (the signal-to-noise ratio or SNR). This can be ameliorated by introducing a stronger primary magnetic field, B_0 , or by acquiring more signals over time. *In-vivo* scan times are generally limited by how long a person can remain still (10 mins is the general upper bound, though multiple scans may be combined). Magnetic field strengths are limited by hardware, with recent innovations including stronger field strengths (Ladd *et al.*, 2018). Thus, holding hardware constant, there is a general trade-off between resolution, SNR, and scan time. Note that resolution scales to the third exponent with scan time (for example, going from 1mm^3 to 0.5mm^3 would require $2^3=8$ times longer to cover the same volume). Additionally, these smaller voxels will now contain $\frac{1}{8}$ the original number of protons, and thus roughly $\frac{1}{8}$ the SNR. Averaging data can recover up to the square root of this SNR, and so to go from 1mm^3 to 0.5mm^3 with equivalent SNR would require $8^2=64$ times more scan time in total. Many clever heuristics and optimizations have also been developed to accelerate acquisition or conversely improve resolution or contrast, such as motion tracking, parallel 2D slice acquisition, efficient or sparse k-space sampling, super-resolution sampling, and many others (see Tsao and Kozerke, 2012; Ye, 2019; and, specifically in the hippocampus, Winterburn *et al.*, 2013). Future work will likely yield further improvements in the image SNR, resolution, and scan time trade-off seen in MRI.

1.4.2 MRI of the hippocampus

The hippocampus can easily be seen in its entirety in MRI, but the laminar and cytoarchitectonic features that define its subfields are not readily available with the sensitivity of human *in-vivo* MRI. Nonetheless, recent advances in MR acquisition methods and the advent of higher magnetic field strength scanners (e.g. 7-Tesla MRI) have greatly improved available image features. In the 1990s and early 2000s, focused structural scans could seldom be acquired below 1mm^3 , and there were often no discernable features within the hippocampus and so it would be given one single

volumetric label (e.g. Filipek *et al.*, 1994; Frazier *et al.*, 2005; and many others). Concurrently, mounting histological evidence further implicated subfield-level pathology in many diseases. This was and still is most typically examined in coronal slices through the hippocampal body, where the arrangement of subfields is simplest since such slices are often roughly perpendicular to the anterior-posterior curvature of the hippocampus (see Figure 3). For these reasons, and in keeping with coronal histological reference materials, initial protocols for estimating the subfields in MRI were optimized for high in-plane coronal resolution with relatively thick slices, typically on the order of $0.4 \times 0.4 \times 2.0 \text{mm}^3$ (Mueller *et al.*, 2008; Yushkevich, Amaral, *et al.*, 2015; Berron *et al.*, 2017; and many others). Scans of this nature have been the primary focus of international harmonization efforts to standardize such protocols (Wisse *et al.*, 2017; Olsen *et al.*, 2019).

Isotropic voxel acquisitions have also been used in some protocols (Wisse *et al.*, 2012; Winterburn *et al.*, 2013), which is advantageous in regions of the hippocampus that show more complex folding, or curve in directions that are not coronal, such as the hippocampal head and tail. This is especially critical as these regions show the most complex anterior-posterior curvature and digitations, though such features are sometimes present in the hippocampal body as well. Anisotropic acquisition methods as described above may cause blurring between digitations, even in the hippocampal body where anterior-posterior curvature is close to coronal oblique. Indeed, the prevalence of anisotropic acquisition protocols may have obfuscated the broad inter-individual variability seen in digitations in the hippocampal body. It should be noted that efficient 3D acquisition methods, as discussed above, often require full brain coverage to eliminate aliasing, which could be an advantage for whole-brain studies but a limitation for studies interested specifically in the hippocampus.

Another common practice in neuroimaging of the hippocampus is using T2-weighted (T2w) acquisition methods, which shows generally inverted contrast to the more common T1w imaging. T2w imaging is particularly advantageous for detecting the SRLM (e.g. Mueller *et al.*, 2008) because of its high myelin content. In principle the SRLM could also be seen in T1w images, however, most researchers have still found T2w to be clearer

for visualizing this structure, possibly because of a contribution from blood vessels inside the hippocampal sulcus which appear similarly dark in T2w images and have less of a contribution in T1w images (see discussion from Kerchner *et al.*, 2012). Thus, it should be noted that the structure generally referred to as high myelin SRLM based on T2w imaging may also contain additional structures like blood vessels. For these reasons, MRI data used in the projects presented in this thesis will consist of isotropic (or nearly isotropic) T2w data.

1.4.3 Manual subfield delineation

One critical intra-hippocampal image feature that is available with sub-millimetric scanning is the high myelin SRLM described above (Kerchner *et al.*, 2010). The SRLM strata surround the hippocampal sulcus which separates the innermost subfield DG from the other subfields curled around it. This feature has been an anchor point in all of the *in-vivo* subfield segmentation protocols referenced here or summarized in (Yushkevich, Amaral, *et al.*, 2015). Other subfields are often defined manually relative to this structure. Here is an example of how a manual subfield segmentation protocol might apply subfield boundaries relative to this structure: the SRLM may make up the superior boundary of subfield CA1, whereas the lateral boundary may be defined perpendicular to the midpoint between a line connecting the most lateral part of the SRLM and the most medial part of the opening of the hippocampal sulcus (Berron *et al.*, 2017). However, it should be noted that over 20 protocols are currently in use across different labs, and can vary in the precise borders, landmarks, labels, and anterior-posterior coverage of the hippocampus (Yushkevich, Amaral, *et al.*, 2015). This has generated confusion and conflicting accounts in the literature, which has been a major driving force for international efforts to harmonize analysis protocols (Wisse *et al.*, 2017).

Typically, histological data would be considered the ground truth in such segmentation protocols, and a manual segmentation of MRI based on this reference would be considered the gold standard. This approach works reasonably well in the body of the hippocampus but has serious limitations in the hippocampal head and tail and does not take full advantage of the 3D nature of MRI. Specifically, alignment of landmarks is highly dependent on the slice angle and distance along the anterior-posterior axis of the

hippocampus (Ding and Van Hoesen, 2015; Cai *et al.*, 2019; de Flores *et al.*, 2019; DeKraker *et al.*, 2020). Inter-individual variability in the anterior-posterior curvature of both the whole hippocampus and its digitations further exacerbate this problem (Ding and Van Hoesen, 2015; Cai *et al.*, 2019; DeKraker *et al.*, 2020). We have discussed these issues and possible solutions in a recent Opinion paper (DeKraker *et al.*, under review). Some of those discussion points are highlighted in the remainder of this section. Note that this paper references published work included in Chapters 2 and 3.

Recent studies utilizing 3D histology, ultra-high field MRI, and *ex-vivo* MRI have begun to reveal the complexity and inter-individual variability seen in folding, or digitations, within the hippocampal head (anterior) and tail (posterior). Figure 3 illustrates this complex folding using data from BigBrain 3D histology (Amunts *et al.*, 2013) and published anatomical work on the hippocampus to be discussed in Chapter 3 (DeKraker *et al.*, 2020). Note that the SRLM appears light here as it contains few stained cell bodies and represents the space between subfields. In this example, it can be easily appreciated that the topological relationship between subfields, including their continuity, is typically not preserved in coronal slices through the hippocampal head and tail, and can change drastically between successive slices. Thus, it can be very difficult to match corresponding slices between histological reference materials and *in-vivo* MRI coronal slices.




Figure 3. Illustration of hippocampal subfield topology in traditional and multiplanar histological slices. A) shows traditional hippocampal slices with and without subfields overlaid. Black lines indicate the continuity of subfields from the dentate gyrus to medial temporal lobe neocortex. B) shows the approximate locations of the sagittal slice shown (outlined in orange) in the full brain, and coronal slices (outlined in solid or dashed blue lines) on the left 3D model of the hippocampus. In particular, slices through the hippocampal head (light blue dashed line) and tail (dark blue dashed line) show within-plane discontinuities. C) shows multi-planar resampled slices along the axis of hippocampal curvature with and without subfields overlaid. The colour of each image outline corresponds to the location of its sampling on the right 3D hippocampal model in B). The continuity of subfields is again shown by a black line and is consistent in all slices in C).

Early histological descriptions of the internal architecture of the hippocampus emphasize its topological continuity despite its complex 3D folding (Nieuwenhuys, Voogd and van Huijzen, 1981; Duvernoy, Cattin and Risold, 2013). More recently, it was shown that reslicing the hippocampus perpendicular to this curvature can simplify the arrangement of subfields in each slice (Gross *et al.*, 2020). This could allow for the application of manually defined labels in each slice to a much simpler configuration of tissue, which

more often resembles the classic curled ‘C’ shape seen in coronal slides through the hippocampal body. In this case global 3D curvature is implicitly encoded in the slice angle. In Figure 3C we performed a similar reslicing of the isotropic BigBrain left hippocampus, and additionally included slices through the most antero-medial ‘vertical component of the uncus’, where the orientation of this axis is axial. As expected, the topology is consistent between all slices, and thus we can imagine it would be much easier to define manual or automated rules for segmentation in these planes than in traditional coronal slices. However, a challenge arises due to variability seen between individuals. This can include gross differences in long-axis curvature, as well as smaller or meso-scale differences in folding or digitation. These features are illustrated using 7-Tesla MRI data from (DeKraker *et al.*, 2018) in Figure 4, and examined in more detail in Chapter 2. The hippocampal subfields, as estimated in (DeKraker *et al.*, 2018), shift in concert with this digitation and curvature. This adds a layer of complexity to both coronal and multi-planar manual segmentation that, to our knowledge, has not been addressed in any manual coronal or multi-planar protocol.




Figure 4. Variability in overall hippocampal shape in four representative individuals as measured in MRI. 3D models are shown to the left of each grey box, with a coronal slice through the hippocampal head (outlined in blue) and a sagittal slice (outlined in orange) to the right. Hippocampi were affine registered such that all slices shown here are from equivalent positions. Hippocampi were organized along two arbitrary, qualitative axes: curvature and digitation. Curvature refers to the long-axis, which can appear more ‘J’ shaped with a straight hippocampal body, or more ‘C’ shaped with a curved hippocampal body. This is most easily seen in the

3D model. Digitation refers to the number and amplitude of grey matter folds within the hippocampus, and can be seen on the 3D model, the coronal plane when digitations are in the hippocampal head, and the sagittal plane when digitations are in the hippocampal body. Equivalent coronal planes can show dramatically different subfield organization depending on these morphological differences.

1.4.4 Registration-based subfield delineation

Given the absence of histological cues in MRI, the task of subfield segmentation is essentially one of manual or computational alignment to histological reference materials. Computationally, this entails maximizing the alignment of image features in one image with another. Typically, one image is held constant (the reference image) while the other (the target or floating image) is linearly translated or stretched in 3D (affine transformations, which can also include rotations, shears, and some other operations), or differentially deformed at all locations until some measure of image alignment is optimized (the cost or loss function). One common example of such a loss function might be the mean squared difference of all voxel intensities between the two images. For images with different contrasts, the cross-correlation between voxels of the two images could be used instead of mean squared difference, allowing for structures of different contrasts but similar shapes to still be aligned during optimization. Many implementations exist to perform this optimization. Optimization of deformable transformations are more challenging than affine transformations given the large number of possible solutions (or degrees of freedom). Thus, most methods introduce some form of regularization to this problem. For example, enforcing similar or smooth deformations between neighbouring voxels effectively reduces the problem space. In many implementations, parameters can be adjusted to make such a registration more fluid, with high degrees of freedom, or more elastic, in which case less deformation will be applied to the floating image such that it might retain some of its original form (e.g. see Dale, Fischl and Sereno, 1999; Fischl and Dale, 2000; Jenkinson and Smith, 2001; Fischl *et al.*, 2002; Jenkinson *et al.*, 2002; Avants *et al.*, 2014; Tustison *et al.*, 2014).

Though initially treated with skepticism due to limitations in implementation (Wisse, Biessels and Geerlings, 2014), computational approaches are becoming increasingly

popular in hippocampal subfield delineation. Such methods are advantageous not only because they are time-efficient and reproducible, but in principle they can also utilize 3D image features that would be out-of-plane to a histologist or manual rater of MRI images. Common approaches to perform segmentation involve deformable registration of a given subject's hippocampus to a sufficiently detailed 3D reference material, which is constructed via a combination of densely sampled histology, *ex-vivo* MRI, and/or manual annotations at very high resolution (Iglesias *et al.*, 2015; Yushkevich, Pluta, *et al.*, 2015; Schmidt *et al.*, 2018). Thus, the quality of segmentation depends on the quality of manually annotated references which are also subject to the problems described in Chapter 1.4.3. Recent and ongoing work enhances both the registration process and 3D reference materials used in these methods (Iglesias *et al.*, 2015; Yushkevich, Pluta, *et al.*, 2015; Wisse *et al.*, 2016; Berron *et al.*, 2017). Automatic Segmentation of Hippocampal Subfields (ASHS) (Yushkevich, Pluta, *et al.*, 2015) is a good example that allows for multiple, interchangeable reference materials. Given good quality reference materials, these approaches can often account for anterior-posterior differences in subfield arrangements and inter-individual differences in the gross curvature of the hippocampus, though inter-individual differences in finer scale curvature and digitation are typically not respected with these methods.

At a finer spatial scale, there are inter-individual differences in digitation within the hippocampus that are not well posed for registration-based approaches to segmentation. For example, it is not clear what deformation should be applied to align a hippocampus with three anterior digitations to one with five anterior digitations. Some solutions may involve topological breaks – combining digitations together or else stretching one digitation out over multiple digitations of a reference image, which can lead to major image distortions. This is similar to the problem of registering the neocortex of individuals with variable gyrification or sulcal patterns, a problem which is discussed in Figure 5 below. The fact that these digitations sometimes continue into the more posterior hippocampal body and tail (DeKraker *et al.*, 2018, 2020; Cai *et al.*, 2019) further enlarges the problem space. These morphological features are already apparent under many state-of-the-art MRI acquisition methods, and so it is not clear that resolution and contrast improvements in acquisition will improve the ability of the approaches discussed above

to capture such details. Rather, improving acquisition protocols may even obviate this shortcoming.

Several notable additions to volumetric, computational segmentation approaches have recently been proposed to improve their generalizability across inter-individual differences in curvature and digitation. One approach - manifold learning for atlas propagation - involves the selection, weighting, or step-wise registration to multi-atlas references that are most suitable for the subject at hand (Wolz *et al.*, 2010). For example, in the case of the hippocampus, a reference atlas could be selected with comparable curvature and digitations to the given subject. This approach was taken by (Xie *et al.*, 2017) for resolving inter-subject variability in perirhinal cortex gyrification, and by (Li *et al.*, 2015) for resolving the entire hippocampus, but it has not been used for hippocampal subfields.

1.4.5 Classification-based subfield delineation

Another approach to segmentation is voxel-wise classification using local or global image features, as in Freesurfer's Adaptive Segmentation (ASEG) (Fischl *et al.*, 2002). In short, a list of features is generated for each voxel, including things like intensity, intensity of neighbouring voxels, likely classes of these neighbouring voxels, or absolute position in the brain. These features are then compared to a reference distribution and the most likely class for each voxel is determined, in this case using a Bayesian mathematical framework. Thus, as in registration-based methods, the reference materials and the optimization algorithms are critical in ASEG, however, this approach is not necessarily limited by deformation parameters like elasticity (though in practice such constraints are often still introduced under a classification-based approach). This approach is used in conjunction with 3D deformable registration in the Freesurfer hippocampal subfield pipeline, which performs a bijective registration and ASEG optimization to segment the hippocampal subfields (Van Leemput *et al.*, 2009; Iglesias *et al.*, 2015), and thus may exhibit some limitations in generalizability as discussed above.

1.4.6 Deep learning-based subfield delineation

A newly popular approach to image segmentation is deep learning, as often implemented using artificial neural networks (ANNs). In general, ANNs iteratively refine or optimize a model with a relatively large number of parameters, rather than more traditional approaches that rely iteratively or analytically optimizing a relatively simple (few parameter) model. In ANNs, the model, or network, typically consists of a series of hierarchically connected nodes, or neurons. We can imagine a set of connections that might map any input, encoded by a subset of neurons, to any output, encoded by another subset of neurons. These encoding schemes tend to be relatively simple, such as a one-hot-vector in the case of classification problems, or pixel/voxel intensities in the case of 2D or 3D images. Training an ANN consists of presenting an input and allowing its values to propagate through the network, which is initially given random connection strengths (these are typically referred to as the model's parameters). Once this signal reaches the output nodes, a difference is calculated between the returned values and the "ground truth", or "supervision" signal. This difference, or loss, is then used to update network connections. By iterating over a huge number of input-output pairs, the network eventually leverages general relationships between the input and output datasets (or distributions). The generalizability of such relationships can then be tested on independent datasets, where ANNs show state-of-the-art performance in many tasks. Many training algorithms have been proposed to optimize this mapping of inputs to outputs in a given task context, but the method of backpropagation of error (or simply backprop) is most ubiquitous (Rumelhart, Hinton and Williams, 1986; Goodfellow, Bengio and Courville, 2016). Briefly, this entails calculating not only the loss from the output, but also a local error estimate for each neuron in a hierarchical fashion (i.e. asking "what should this neuron's activity have been in order to generate the correct response for its downstream neurons?"). Though many additional innovations have since been added, this simple learning rule has proved amazingly robust and general, and it now makes up the backbone of nearly all deep learning applications.

U-Net (Chen *et al.*, 2018) is a popular ANN architecture for image segmentation that makes use of recent innovations such as convolutional layers (Krizhevsky, Sutskever and

Hinton, 2017), residual connections (He *et al.*, 2016), and an encoder-decoder framework with a bottleneck layer (Ronneberger, Fischer and Brox, 2015; Badrinarayanan, Kendall and Cipolla, 2017). The utility of an encoder-decoder framework has to do with the relationship between generalizability and compression. For example, if we want to classify images of cats and dogs, each image may consist of one thousand pixels. There is no one pixel value or relationship between pixels that make up a general category of dogs or cats, but by compressing such images down to a relatively small collection of features makes such a distinction much easier. For example, one feature might be any arrangement of pixels into something resembling whiskers, the presence or absence of which could then be used to classify the image as a dog or a cat. An ANN is not provided with or restricted to any specific feature, but it tends to naturally produce features that differentiate the categories of interest. Interestingly, such features are often naturally recognizable by humans (Olah *et al.*, 2018) and tend to correlate with estimates of representational content in the brain (Cichy *et al.*, 2016). This extraction of features, or compression, comprises the ‘encoding’ branch of an encoder-decoder architecture. The ‘decoder’ generally performs a similar decompression of features into an output with corresponding pixels as the input, and the output resembles an input image in that it could have plausibly produced the same set of compressed features. In the context of biomedical imaging, there is no one voxel that makes a brain structure recognizable and so feature extraction, or compression, is also advantageous. For example, using an encoder-decoder framework the input may be an MRI image and the output may be a segmentation of corresponding labels, and the compressed features (or the bottleneck layer) might consist of local and global relationships between pixels.

Convolutional layers in the encoder or deconvolutional layers (also known as transpose convolutions, or dilated convolutions) in the decoder can make networks like U-Net easier to train in practice by reducing the overall number of parameters compared to a fully connected network. This works by reusing a set of relationships between pixels, or filters, across multiple locations in an image. This is similar to traditional image processing where an image may be passed through multiple handcrafted filters to highlight certain image elements, but in convolutional ANNs the nature of the filters is learned from the bottom up during training.

Finally residual and ‘skip’ connections bypass some hierarchical steps of an ANN architecture, which improves U-Net performance in two ways: typically there is a drop-off in performance as ANNs become too deep, but introducing residual connection helps to mitigate this problem (known as the vanishing or exploding gradients problem) (Zaeemzadeh, Rahnavard and Shah, 2020). Secondly, residual connections can provide low-level or perceptually detailed features from the encoder to the decoder without having to pass through the more compressed bottleneck layer, which helps produce more detailed output segmentations.

Some recent protocols apply ANNs, including U-Net, to voxel-wise classification of hippocampal subfields (Shi, Cheng and Liu, 2019; Zhu *et al.*, 2019; Yang *et al.*, 2020). In principle, ANNs could be robust to inter-individual differences in hippocampal curvature and digitation or could leverage subtle intensity differences indicative of the underlying cytoarchitecture. However, in each of the cases cited here, output segmentations closely resemble the manual segmentations that they were trained on, which do not explicitly address intensity differences between subfields or differences in digitation and curvature between subjects. Generally, the use of ANNs in neuroimaging tends to be limited by the availability of detailed and consistent training data. I will explore this class of approaches using more detailed reference materials and in conjunction with a surface-based approach in more detail in Chapter 4.

1.4.7 Surface-based subfield delineation

The logic of surface-based registration is to account for inter-individual differences in folding by projecting to a 2D flat or spherical surface, which removes gyrification, curvature, and digitations but preserves topology. Once these patterns of folding are detected and unfolded, the problem of defining hippocampal subfield boundaries becomes much simpler (i.e. by constraining it from a 3D problem to a 2D problem). Registration and parcellation performed in such an unfolded space can then be projected back to the native space, an approach that has been highly successful in the neocortex (Dale, Fischl and Sereno, 1999; Fischl, Sereno and Dale, 1999). An example that may be familiar to medial temporal lobe researchers is the collateral sulcus, which shows highly variable depth and branching patterns (Huntgeburth and Petrides 2012; Berron *et al.*

2017; Pruessner *et al.* 2002). It is not clear what 3D deformation should be applied to align a subject with a branched or split collateral sulcus to a subject with a single collateral sulcus (green arrows in Figure 5), but a surface-based representation allows an equivalent surface to be fit to both cases, as illustrated in the example shown in Figure 5. Furthermore, the boundaries of parcels or ROIs can be aligned in 2D on such a surface, allowing them to be shifted topologically between individuals without erroneously crossing sulci or gyri (known as topological breaks), which can cause a parcel to become drastically stretched or compressed. Typically 2D neocortical surfaces are aligned according to fixed anchor points (e.g. the edges of the neocortex), and additional features can be used to further refine registrations, for example using the major gyri and sulci evident in highly smoothed gyrification maps (Dale, Fischl, and Sereno 1999; Desikan *et al.* 2006; B. Fischl, Liu, and Dale 2001; B. Fischl, Sereno, and Dale 1999) or intracortical myelin maps and other features (Glasser *et al.* 2016). This is particularly advantageous for parcellation driven based on cytoarchitectonic features (as opposed to parcellation based on connectivity or task involvement), which are generally topologically clustered due to their ontogeny in the developing neural tube (Sperber, 1995; see Figure 1). In principle, this approach can accommodate any pattern of digitation and curvature within the hippocampus, and potentially even abnormal cases like hippocampal malrotations (Lehéricy *et al.*, 1995; Tsai *et al.*, 2016). As such, this type of approach is promising for subfield segmentation, and more generally, for measuring structural features of the hippocampus at a level of detail that requires delineation of individual digitations.




Figure 5. Equivalent coronal slices showing the left collateral sulcus in two Human Connectome Project subjects (Elam and Van Essen, 2015) (minimally preprocessed and parcellated using Freesurfer’s surface-based APARC neocortical parcellation). Labels include hippocampus (orange; ASEG subcortical parcellation method),

parahippocampal cortex (green), fusiform cortex (sage) and inferior temporal cortex (pink).

Several studies have used surface-based methods to ‘flat-map’ the hippocampus (Zeineh *et al.*, 2001; Ekstrom *et al.*, 2009; DeKraker *et al.*, 2018, 2020; Vos de Wael *et al.*, 2018) (Note that DeKraker *et al.*, 2018 and 2020 represent the work described here in Chapters 2 and 3). In its first implementations, Bookheimer and colleagues (Zeineh *et al.*, 2001; Ekstrom *et al.*, 2009) performed subfield segmentation in native volumetric space, before transforming the entire volume of the hippocampus, with its segmentations, to an unfolded space. In this approach, all subfield labels and corresponding MRI voxels are transformed to flattened space based on multidimensional scaling methods similar to those that have been used in the neocortex (Engel, 1997; Lauer *et al.*, 2018). Other groups have employed similar combinations of manual and surface-based methods, but with flat-mapping applied at the level of individual subfields which are then stitched together to make up a single surface after the fact (Vos de Wael *et al.*, 2018). Chapters 2 and 3 show the development and applications of an approach that starts with unfolding of the entire hippocampal volume and subsequently performs segmentation in this unfolded space. All of these methods (Zeineh *et al.*, 2001; Ekstrom *et al.*, 2009; DeKraker *et al.*, 2018, 2020; Vos de Wael *et al.*, 2018) allow for measurement of structural hippocampal features, such as thickness, by virtue of providing a dimension that can be estimated perpendicular to the axis of hippocampal topology. However, only approaches that start with unfolding (DeKraker *et al.*, 2018, 2020) explicitly leverage topology to register to a standardized space and perform subfield delineation in a topologically informed and contiguous manner.

1.5 Goals of this thesis and overview of projects

Broadly, the goal of this thesis was to gain as much structural information as possible from neuroimaging of the hippocampus. To that end, this thesis aimed to develop and explore a topological (or surface-based) coordinate framework for analyzing the hippocampus in 3D. This framework should encompass inter-individual differences in intra-hippocampal folding, a major methodological challenge in the field. Ideally, this would enable topological alignment between subjects with different native 3D

hippocampal folding configurations. Under this framework, hippocampal subfields should be definable in an unfolded 2D space, simplifying the problem from its original 3D form where the spatial relationship between subfields can be highly complex, especially in the hippocampal head and tail. Additionally, based on the literature, one would expect that subfields should be fully contiguous when unfolded. Such a framework would go beyond extant surface-based hippocampal analysis methods by leveraging inferences about hippocampal topology or folding as a basis for alignment and subfield definition, as opposed to other methods that reconstruct or align surfaces from predefined or heuristically estimated subfields in native 3D space. The methods developed here should also apply a higher level of detail than extant methods, ideally delineating even small folds (or digitations) within the hippocampus. This would improve measures such as thickness, which might otherwise erroneously cross hippocampal folds. Finally, with such a framework it should be possible to do the following: i) sample quantitative MRI with minimal partial voluming effects from surrounding or intra-hippocampal white matter, ii) given sufficient resolution and contrast, extract laminar features of hippocampal tissue, and iii) quantify other morphological characteristics of hippocampal tissue, such as the extent of hippocampal folding, or gyrification. Overall, these methods should flexibly encompass inter-individual variability in hippocampal folding, constrain possible subfield definitions, and improve the sensitivity and specificity of hippocampal structural measures.

Chapter 2 aimed to develop a computational method to unfold hippocampal grey matter, with a focus on the hippocampal head where complexity is highest due to medial curving of the structure and the variable presence of digitations. In particular, we aimed to leverage the visibility of a critical image feature composed of the hippocampal sulcus and strata radiatum, lacunosum, and moleculare, (SRLM) in high resolution neuroimaging to facilitate the separation of different hippocampal folds or digitations. This structure was segmented in high-resolution, T2-weighted 7-Tesla MRI data from 12 healthy participants and one surgical patient with epilepsy whose resected hippocampal tissue was used for histological validation. We then employed user-guided semi-automated techniques to detect and subsequently unfold the surrounding hippocampal grey matter. This unfolding was performed by solving Laplace's equation in three dimensions of

interest (long-axis, proximal-distal, and laminar). The resulting ‘unfolded space’ provided an intuitive way of mapping the hippocampal subfields in 2D space (long-axis and proximal-distal), such that similar subfield borders can be applied in the head, body, and tail of the hippocampus independently of its folding. This unfolded space was additionally employed to map intracortical myelin and thickness in relation to subfield borders, which revealed intracortical myelin differences that closely followed the subfield borders used here. Examination of a histological resected tissue sample from a patient with epilepsy revealed ways in which our unfolded coordinate system has biological validity. This examination also showed that subfield segmentations applied in this space are able to capture features not seen in extant fully manual tracing protocols.

Chapter 3, we asked whether the unfolded coordinate framework we had previously developed might be leveraged to extract features there are sufficient to derive the known hippocampal subfields. For this, we analyzed a histological dataset with unprecedented 3D coverage, BigBrain. From this data we imposed the computational unfolding framework developed in Chapter 2. We adapted neocortical feature extraction techniques to our coordinate framework to map the hippocampus with respect to 10 laminar and 5 morphological features. Unsupervised clustering of these features revealed subdivisions that closely resemble gold standard manual subfield segmentations in the native 3D histology. Critically, we also showed that morphological features alone are sufficient to derive most hippocampal subfield boundaries, despite the fact that hippocampal subfields are traditionally defined by laminar composition. This offers promise for *in-vivo* neuroimaging where laminar features are typically not available. Moreover, some features showed differences within subfields along the hippocampal longitudinal axis. These findings highlight new characteristics of internal hippocampal structure and offer new avenues for its characterization with *in-vivo* neuroimaging.

Chapter 4 aimed to generalize the methods developed throughout this thesis to new datasets in a way that can readily be taken up by new labs without manual segmentation, which is costly in both time and expertise required. Specifically, the unfolding framework developed in Chapters 2 and 3 is initialized by classification of the different intra-hippocampal tissues, such as grey matter, SRLM, cysts, and other surrounding structures.

These tissue classes must be delineated with sufficient detail to differentiate distinct folds within the hippocampus. Similar to subfield segmentation, this problem is challenging to address using traditional medical image processing techniques due to broad inter-individual variability in hippocampal folding configurations. Thus, we employed a U-Net deep learning architecture. We overcame limitations in the amount of available training data by pooling data across previous studies and by applying incremental learning over the open source Human Connectome Project dataset. Final hippocampal tissue segmentations on left out test data showed high Dice overlap with manual segmentations. Following post-processing to correct for segmentation errors (specifically topological breaks which distort subsequent unfolding), seamless integration was seen with the unfolding coordinate framework, with similar patterns of gyrification, quantitative T2w, and thickness as seen in previous work. We applied histologically defined subfield boundaries from 3D histology as examined in Chapter 3 to all subjects in unfolded space, and propagated labels back to each subject's native space hippocampal folding configuration. We compared these subfield segmentations to other prominent automated methods and found that our subfield segmentations generally compare favourably to other methods in terms of subfield detail, coverage of the full hippocampus, and definition of digitations.

Chapter 2

2 Unfolding Coordinate Framework

This Chapter focuses on the methodological details of the proposed standardized topological coordinate framework for the hippocampus, which is now published in the journal *NeuroImage* with the title ‘Unfolding the hippocampus: An intrinsic coordinate system for subfield segmentations and quantitative mapping’ (DeKraker *et al.*, 2018). This framework is not inherently specific to any imaging modality, but in the current study data were acquired from 7-Tesla structural MRI (T2w; 0.3mm³ final resolution). The hippocampus, including intra-hippocampal high myelin strata and several surrounding structures, were manually delineated. Gradients were then imposed over hippocampal grey matter along three dimensions: anterior-posterior, proximal-distal, and inner-outer. These gradients were defined by solving Laplace’s equation with boundary conditions where the hippocampus borders neighbouring structures. For example, the anterior-posterior gradient boundaries were defined by the border between hippocampal-amygdalar transition area and the antero-medial hippocampal head and the border between the indusium griseum and the postero-medial hippocampal tail. Indexing hippocampi according to these gradients provided a method for implicit topological registration to a standardized or ‘unfolded’ space. The utility of this method was demonstrated in two ways. First, quantitative features (intracortical myelin and cortical thickness) were mapped in unfolded space for all subjects. Averaging across participants in this topologically aligned space gave rise to emergent patterns which qualitatively match observations from *ex-vivo* studies of the hippocampus. Secondly, each hippocampal subfield was mapped in unfolded space, where their contiguity can be easily observed, and a surprisingly high degree of inter-subject consistency was seen that is not evident in native 3D space. These subfield definitions were validated in a patient MRI and corresponding *ex-vivo* histology sample.

2.1 Introduction

Researchers often distinguish the hippocampus from neocortex but the hippocampus, in fact, also has a cortical composition sometimes referred to as archicortex due to its wide

evolutionary preservation (e.g. Duvernoy *et al.*, 2013). Like the neocortex, the hippocampus shows variable gyrification, often referred to as digitations or pseudo-digitations in the anterior hippocampal head and more posterior body/tail, respectively. This variability creates major challenges for cross-participant alignment and segmentation. This is particularly of interest given the recent controversy over segmentation of the hippocampus into subfields in MR data, which are not sensitive to most cytoarchitectonic features that define the subfields (for an overview of this controversy see Yushkevich *et al.*, 2015 and harmonization efforts by Wisse *et al.*, 2017).

Though present in the rest of the hippocampus, digitations are most prominent in the hippocampal head. This structural feature is a significant challenge for subfield segmentation protocols and as such most protocols do not segment this region, or do not honour its complex and variable structure (see Yushkevich *et al.*, 2015). Ding and Van Hoesen (2015) recently provided detailed descriptions of the hippocampal head including three different morphologies (2, 3, or 4 digitations). However, there are observed cases with even more digitations that continue through the hippocampal body (see also Gao and Ver Hoef, 2016) and cases with differences in the amount of medial curvature of the uncus. Dalton *et al.* (2017) and Berron *et al.* (2017) have recently published protocols leveraging Ding & Van Hoesen's descriptions. These protocols collapse across different morphologies and deal primarily with one canonical case. This may produce results that are close to the ground truth under different morphologies as well. However, differences in folding will cause a topological shift and so each subfield border should shift in turn. Thus, attempting to impose borders without considering topology creates challenges in subjects with different degrees of folding, or different rotations or positions within the medial-temporal lobe (e.g. varying degrees of dysplasia), similar to the challenge of aligning the neocortex in participants with variable gyrification.

In the neocortex, the challenge of inter-subject alignment in cases of variable gyrification have been largely overcome by using topology-preserving surface-based alignment (Dale *et al.*, 1999, Fischl *et al.*, 1999a, Fischl *et al.*, 1999b, Fischl *et al.*, 2001), which has led to the development of powerful methods for parcellation (e.g. Glasser *et al.*, 2016). These types of methods have not yet been applied to the archicortex of the hippocampus.

However, several studies reported by Bookheimer and colleagues have implemented a technique that is similar but used primarily for visualizing results rather than as an analysis technique (see Ekstrom *et al.*, 2009, Zeineh *et al.*, 2003; and also in 7-Tesla MRI Suthana *et al.*, 2015). Under their protocol, delineation of medial temporal neocortex and hippocampus is performed in the subject's native space and conformal mapping is used to flatten this tissue, such that results can be viewed in a single plane. However, this protocol does not make use of some of the advantageous features used in neocortical surface-based analysis: it does not use a standardized set of coordinate points and segmentation is performed based on a geometric landmarks in native space (i.e. prior to any surface-based alignment). Although the folded topology in the hippocampal body is captured, the digitations and medial curvature of the hippocampal head and tail do not appear to be separately delineated – instead they are labelled using a similar coronal scheme as the hippocampal body. Consequently, topology is not fully preserved in these areas under this protocol.

2.1.1 Goals of the current study

The current study aimed to examine the topological structure and ontogeny of the hippocampus, as discussed above, in order to develop a two-dimensional coordinate system for alignment and segmentation of variably folded hippocampi across individuals, similar to surface-based alignment methods used in the neocortex. Specific structural features we identified and aimed to account for are the medial folding forming the classic hippocampal C-shape (or inverse C-shape depending on hemisphere and orientation), long-axis and uncus curvature, digitations, and inter-individual variability in each of these features. After tracing each of these features in 7-Tesla T2-weighted MR images, we applied the Laplace equation to divide hippocampal archicortex into a set of standardized long-axis and proximal-distal coordinates using anatomically motivated boundary structures that are topologically continuous with the hippocampus. We applied a segmentation of the hippocampus based on the histological samples used by (Ding and Van Hoesen, 2015) under the framework of this two dimensional, topology-preserving coordinate space, which we then validated by comparison to quantitative MR measures of intracortical myelin and thickness, as well by direct comparison to a surgically resected

tissue sample from a patient with epilepsy (i.e. comparison of preoperative segmentation to postoperative histology).

2.1.2 Critical structural features we aimed to accommodate

As discussed above, during development the hippocampus originates from a single flat tissue, which in addition to its long-axis curvature, also folds medially upon itself forming a C-shape while differentiating into the various subfields. This developmental characteristic has several interesting consequences for the structure of the adult hippocampus: all subfields make up adjacent segments of a contiguous tissue (though the dentate gyrus makes up a distinct tissue but keeps a consistent position at the distal edge of the CA fields). The sulcus, or ‘crease’, around which this folding occurs can be visualized in histology as the hippocampal sulcus and surrounding high myelin laminae SRLM. In the current study we aimed to capture the SRLM in the hippocampal head and tail as well, which we then critically leveraged to differentiate the folds of the entire hippocampus, preserving its topology.

Recent histological evidence from (Ding and Van Hoesen, 2015) offers a new morphological characterization of the hippocampal head, which we also aimed to respect in our unfolded coordinate space. A main finding in this characterization was the documentation of considerable inter-individual differences in digitations (i.e. folding, similar to the gyrification of neocortex) in the hippocampal head, varying from 2 to 4 digitations, with additional pseudo-digitations sometimes found along the lateral and inferior sides of the hippocampal body and tail. Ding & Van Hoesen also delineated the subfields in detail in the uncus - a part of the hippocampal head that curves medially, and then superiorly (see Figure 6B). In line with (Duvernoy, Cattin and Risold, 2013), (Ding and Van Hoesen, 2015) showed that all subfields of the hippocampus contiguously follow this curvature through the hippocampal head and have their natural anterior termination not in the absolute anterior tip of the hippocampus, but rather in the more medial and posterior vertical component of the uncus (see Figure 6C). As the subfields curve into the uncus, their borders also shift such that the subiculum and CA1 move from the inferior side to the lateral, anterior, and finally superior side. A detailed segmentation

of this region must capture each of these features, and here we aim to provide a tool with sufficient precision and documented validity to index these structural complexities.

2.2 Methods and data

2.2.1 Study participants

For MRI data acquisition, healthy participants were recruited from Western University, London, Canada ($n = 12$; 6 females; ages 20–35, mean age 27.6). This study was conducted with Western's Health Science Research Ethics Board approval, and informed consent was obtained from each participant prior to participation.

2.2.2 MRI acquisition

Imaging was performed using a 7-Tesla neuroimaging optimized MRI scanner (Agilent, Santa Clara, CA, USA/Siemens, Erlangen, Germany) and employing a 23-channel transmit-receive head coil array, which was constructed in-house. Four T2-weighted turbo spin echo (TSE) 3D (3D sagittal, matrix: 260×366 , 266 slices, 0.6mm^3 isotropic, ~ 8.5 mins per scan) images were acquired from each participant. All images were acquired in sagittal rather than coronal oblique orientation for optimal whole brain coverage, given that these data were also used for other, whole brain studies. The use of isotropic acquisition differs from most hippocampal imaging protocols that acquire thick coronal slices oblique to the hippocampus to maximize in-plane resolution. However, these protocols limit the visibility of structures that run perpendicular to the long-axis of the hippocampus, including most of the hippocampal head and tail. By using isotropic voxels, we were able to capture small features such as the hippocampal SRLM in high detail, throughout the entire hippocampus. A T1-weighted MPRAGE (3D sagittal, matrix: 256×512 , 230 slices, 0.75mm^3 isotropic) image was also collected for all participants.

2.2.3 Preprocessing

All scans were processed as follows: the first T2-weighted image (scan 1) was upsampled to 0.3mm^3 isovoxels using cubic spline interpolation; subsequently, scans 2, 3, and 4 were rigidly registered to scan 1 using FSL FLIRT registration (Jenkinson, 2002, Tofts,

2005). All four scans were then averaged together to produce a single, 0.3mm^3 isovoxel, high-contrast volume, similar to Winterburn *et al.* (2013). This volume was reoriented to an oblique orientation, with coronal slices perpendicular to the long-axis of the hippocampus, by rigid registration to an average template in coronal oblique orientation (as determined by visual alignment of the long-axis of the template hippocampus to the anterior-posterior axis). T1-weighted scans were registered to this high-contrast coronal oblique T2-weighted volume using rigid registration as described above.

2.2.4 Detection and labelling of the SRLM and hippocampal grey matter

Under our isotropic MR acquisition protocol, the SRLM was visible in the entire long-axis of the hippocampus, including the digitations and uncus of the hippocampal head. Representative slices from the hippocampal head and body and a 3D model reconstruction can be seen in Figure 12. All manual tracing was performed in ITK-SNAP 3.4 (Yushkevich *et al.*, 2006), and the built-in ‘Snake’ tool was also used to facilitate tracing. The detailed protocol for tracing, ‘feathering’, dilation, and manual adjustments can be found in Appendix 1.




Figure 6. Illustration of SRLM (green) and hippocampal grey matter (red) labelling.

A) 3D model of the SRLM label in the center, with representative coronal slices

from the head and body on the sides, at the positions indicated by the dotted lines. B) Same views as above but depicting the SRLM label after spherical dilation. C) Same views as above showing the combined SRLM label and grey matter label, after manual adjustments to the grey matter label.

Much of the morphology of the hippocampus is systematically related to the SRLM. This can be observed in our dataset (e.g. Figure 12), where models of the SRLM capture the same digitation and curvature structures as models of hippocampal grey matter. This characterization also agrees with anatomical descriptions wherein all laminae of the CA fields (including SRLM), and even the underlying dentate gyrus, follow the hippocampus' digitated structure (see Figure 11; Duvernoy, Cattin and Risold, 2013; Ding and Van Hoesen, 2015). The various subfields of the hippocampus surround the SRLM, and, thus, we made use of this proximity to initialize grey matter segmentation with the SRLM segmentation. We segmented the grey matter of the hippocampus initially by dilation of the SRLM label using an active contour evolution of the SRLM, followed by manual correction. We used ITK-SNAP's Snake tool (Yushkevich *et al.*, 2006), which evolves a seed region in 3D to fill a structure of interest. In our case, we initialized the evolution using the SRLM and first applied no constraints on the evolution, which resulted in uniform, spherical dilation. The amount of dilation was determined by visually inspecting whether the outer borders of hippocampal grey matter had been reached and varied slightly between traces depending on the available image information. Evolution constrained by edge attraction (with parameters defined by the user based on image quality) and manual adjustments were then applied. We did not trace the SRLM along the superior side of the subiculum (the most medial, less folded extension of the hippocampus) as it was not consistently visible. Accordingly, grey matter in this region was not labelled as part of the dilation of the SRLM label but had to be labelled manually by the rater, using a spherical paintbrush. This was also the case in the most medial, vertical component of the uncus where the SRLM was often not visible. Further manual adjustments included the removal of grey matter label from the CSF on the medial side of the dentate gyrus, and minor changes throughout to ensure all grey matter was labelled as such. Because errors in manual segmentation can produce distortions in the next step of hippocampal grey matter unfolding, the unfolding results of each hippocampus were

visually examined by the raters to ensure their labelling followed the rules outlined in Appendix 1.

To assess how reliably the SRLM could be segmented in our high-resolution images, we repeated the segmentation with an additional trained rater, and calculated the spatial overlap between these segmentations using the Dice similarity index (DSI). DSI represents the proportion of overlapping voxels in two segmentation labels over the mean number of voxels per label. It can vary from 0 to 1, with values close to 1 denoting high overlap (Dice, 1945; Sørensen, 1948).

2.2.5 Manual subfield segmentation

Before unfolding of hippocampal grey matter, we performed manual segmentation of this tissue into subfields in a set of 10 hippocampi (5 participants \times 2 hemispheres) with varying numbers of digitations and varying curvature in the uncus. This was done by carefully matching coronal views in our MR images with the closest corresponding histological segmentations in the hippocampal head provided by (Ding and Van Hoesen, 2015). In the hippocampal body and tail, segmentations were performed based on the descriptions of (Duvernoy, Cattin and Risold, 2013). Representative slices of these segmentations can be seen in Figure 12, and additional examples as well as qualitative descriptions can be found in Appendix 2. Note that it was not our intention to develop a manual segmentation protocol in this paper; we simply aimed to determine whether the trajectories of the hippocampal subfields in our unfolded coordinate space could be captured in a way that respects the recently elucidated complexity in the hippocampal head. Thus, we assessed the reproducibility of these segmentations in only a small sample (four left and four right) using DSI scores as above, to ensure values were at least comparable with previous reports (see results in Figure 10).




Figure 7. Representative slices showing manual subfield labelling from hippocampal head (top) and body (tail). Left shows reference materials from (Ding and Van Hoesen, 2015), middle shows manually traces SRLM (green), hippocampal grey matter (red) and cysts (yellow), right shows manually delimited subiculum (brown), CA1(blue), CA2 (white), CA3 (pink), and dentate gyrus (cyan). See Appendix 2 for further details. Images were adapted with permission from (Ding and Van Hoesen, 2015).

2.2.6 Unfolding of hippocampal grey matter

In the neocortex, 3D computational tools such as Laplace's equation have been used to precisely and flexibly calculate neocortical thickness (e.g. Jones, Buchbinder and Aharon, 2000; Sowell, 2004). In principle, the Laplace equation, $\nabla^2\phi = 0$, defines a potential field (ϕ) whose values change based on their distance from two boundary surfaces. The solution is twice differentiable ($\nabla^2 = 0$), which guarantees a level of smoothness that is appropriate for brain anatomy. In studies of neocortical thickness, the potential field spans the neocortical grey matter, while the boundaries are the white matter and pial surfaces. Thickness is then computed by generating streamlines across the resulting potential field gradient.

We reasoned that Laplace's equation may also be used for unfolding of hippocampal grey matter, not only to determine thickness along the laminar dimension, as above, but also to

compute potential field gradients along the long-axis and proximal-distal dimensions. To do so, it is critical to employ multiple sets of boundary conditions, sometimes referred to as ‘source’ and ‘sink’. For example, unfolding along the long-axis dimension makes use of anatomically motivated boundaries at the anterior (source) and posterior (sink) ends of the hippocampus. The potential field in between is defined over all grey matter and increases smoothly from source to sink. We thus solved Laplace's for three different equations, $\nabla^2\phi_{\text{long-axis}} = 0$, $\nabla^2\phi_{\text{proximal-distal}} = 0$, $\nabla^2\phi_{\text{laminar}} = 0$, to determine a different potential field for each hippocampal dimension. The domain was identical for each dimension, i.e., the hippocampal grey matter, but the boundary conditions were distinct in each of them, defined as anatomical landmarks along at the edges of the hippocampal tissue.

An iterative finite-differences approach was used to obtain the solution for each Laplace equation, while employing a 26-neighbour average to compute the updated potential field and terminating when the potential field change is below a specified threshold (sum of changes $< 0.001\%$ of total volume). It is important to note that the SRLM voxels were not included in the grey-matter domain of Laplace's equation. This protocol feature effectively provides a barrier such that the streamlines follow a geodesic (i.e. along the outer surface) path along the hippocampus and do not introduce short-circuits. Solving of the Laplace equations was performed in MATLAB (code available at <https://github.com/jordandekraker/HippUnfolding>).

Long-axis dimension boundaries

As discussed in Chapter 2.1.2, each of the subfields has its natural anterior terminus in the vertical component of the uncus (Ding and Van Hoesen, 2015). Hippocampal grey matter in this area borders the grey matter of the amygdala, making up an area that is typically referred to the hippocampal-amygdalar transition area (HATA) (Ding and Van Hoesen, 2015). At the tail of the hippocampus, a structure named the indusium griseum (which is actually a vestigial extension of the dentate gyrus) extends medially and posteriorly from the hippocampus and then curves upward and anteriorly along the midline of the brain, before merging with the cingulate cortex (Duvernoy *et al.*, 2013).

The HATA and indusium griseum, thus, make up two visible structures that correspond to the natural anterior and posterior termini of each of the hippocampal subfields. We manually traced these structures only where they border hippocampal grey matter (see Appendix 1 for details) and used them as source and sink regions in Laplace's equation (see Figure 8A for illustration).




Figure 8. Illustration of Laplacian unfolding along the long-axis, proximal-distal, and laminar dimensions in A), B), and C), respectively. The upper left inset image in A) shows a 3D model of the SRLM (green) and grey matter (red) labels, with the HATA (pink) and indusium griseum (grey) to be used as boundaries for Laplace's equation. The lower left image in A) shows arbitrarily coloured bins within the resulting potential field gradient. To the right is the same model as the lower left but

showing cross sections from the body (top) and head (lower three), depicting in particular the uncus (lower two) and vertical component of the uncus (bottom). The locations of these cross sections are shown by the black dotted lines (1–3) and the eye figure (4). B) shows the same views of the same hippocampus but using the inner dentate gyrus (turquoise) and medial temporal lobe cortex border (white) (upper left insert in B), as boundaries for Laplace's equation for the proximal-distal dimension. C) shows the same views of the same hippocampus but using the SRLM (seen in green under semi-transparent red grey matter) and outer hippocampal borders (upper left insert in C) as boundaries for Laplace's equation for the laminar dimension. White dotted lines in C) (right) show the true laminar structure of CA4 and DG, which is not respected in our laminar potential field gradient.

Proximal-distal dimension boundaries

We defined the proximal border as the point at which the subiculum, the most proximal subfield, contacts the grey matter of neighbouring medial-temporal lobe neocortex (see Appendix 1 for details). To index the full extent of hippocampal grey matter, its distal border can be defined as the granule cell layer of the dentate gyrus (i.e. the part of dentate gyrus which most closely borders the SRLM). We employed a custom approach to detect this tissue: within each previously computed long-axis bin we applied volumetric fast marching (Sethian, 1996) along hippocampal grey matter, starting at the border with surrounding temporal lobe cortex, and estimated the dentate gyrus as being the most distal 12% of this distance (determined experimentally). To index only the innermost granule cell layer of this tissue, we dilated the SRLM by a single voxel (8 nearest neighbours) over this rough dentate gyrus approximation. The result included only the most distal portions of the dentate gyrus, corresponding roughly to the granule cell layer.

An additional challenge in the proximal-distal unfolding of the hippocampus lies in the vertical component of the uncus. Here, hippocampal grey matter, unlike in the rest of the hippocampus, does not follow the classic folded C-shape, and instead flattens out (Figure 8 right panels 3 and 4). The dentate gyrus continues on the medial edge of this region, but it is not separated by the other subfields with any visible SRLM at the current image

resolution. We thus defined the dentate gyrus' location manually for this region (Figure 8B, third and fourth right panels, turquoise bin).

Laminar dimension boundaries

We defined the sink for Laplace's equation as the outermost surface of the hippocampus, and the source as the SRLM. However, as mentioned in Chapter 2.2.1, the subiculum and vertical component of the uncus do not border the SRLM. Therefore, these labels were artificially extended over these regions. For the subiculum, this was performed computationally by dilating the SRLM label along the surface of the subiculum until the most medial point was reached in each coronal slice. For the vertical component of the uncus this label was created manually. These artificially extended labels were used as the source in Laplace's equation (Figure 8C).

2.2.7 Subfield borders in unfolded coordinate space

The long-axis and proximal-distal potential field gradients together make up a 2D coordinate system that can be used for indexing columns of hippocampal grey matter. Using this 'unfolded coordinate space', the location of each subfield border can easily be indexed. We used the manual subfield segmentations performed on 10 hippocampi (see Chapter 2.2.5) to generate a subfield atlas in the unfolded coordinate space. That is, for each manually segmented hippocampus, we identified the long-axis and proximal-distal coordinates that correspond to each of the subfield borders. We then averaged these borders together at each long-axis point and plotted the labelled data in the 2D unfolded coordinate space (see Figure 11).

Given the low variability of the subfield borders in unfolded space, we then applied our Laplacian unfolding to the remaining hippocampi and, rather than performing subfield segmentation manually, we applied the group-averaged borders from Figure 11. That is, for each long-axis and proximal-distal coordinate in each hippocampus, we assigned the corresponding label from Figure 11. We then assessed the overlap of the 10 manually segmented hippocampi to their unfolded group-average border segmented counterparts using Dice Similarity Indices. To avoid bias, we used a leave-one-out approach wherein a

given participant's manual segmentations (both left and right) were not included in unfolded group-averaged borders.

2.2.8 Quantitative unfolded tissue properties

Properties such as intracortical myelin content and cortical thickness have been shown to be useful for parcellation of the neocortex into functional subregions (e.g. Glasser and Van Essen, 2011; Glasser *et al.*, 2014). The ratio of T1-weighted over T2-weighted values produces a map that is correlated with quantitative R1 and is used as a surrogate measure of intracortical myelin (Glasser and Van Essen, 2011). We estimated intracortical myelin in this way and estimated cortical thickness by fitting streamlines to the laminar potential field gradients of all hippocampi. We then plotted these values across unfolded coordinate space. To avoid confounds from partial voluming, we mapped the myelin contents of hippocampal tissue from only the middle 25–75%, as determined by our laminar Laplacian field, corresponding to approximately 2 voxels at the current resolution (similar results were obtained when the superficial and deep laminae were included as well). To illustrate how these values map onto the native 3D space of the hippocampus, we generated 3D hippocampal models with surface colouring that corresponds to the underlying myelin estimates from the group average in two hippocampi (one highly digitated and one less digitated exemplar; see Figure 11).

2.2.9 Histological validation

One temporal lobe epilepsy patient with left mesial temporal sclerosis (age 34; male) underwent preoperative 7-Tesla scanning and then went on to receive a left anterior temporal lobectomy, with inclusion of the amygdala and hippocampus, as part of their standard of care. The surgically resected tissue underwent a standardized protocol involving overnight scanning in an ultra-high field *ex-vivo* 9.4-Tesla MRI, agar embedding, and cutting into blocks 4.4mm apart for paraffin embedding and histological sectioning. Staining with H&E, Neu-N, GFAP, and Luxol fast blue was performed, and slides were digitized at 0.5 micron/pixel resolution. The subfields were manually annotated on the Neu-N histology images by Kayla Ferko (K.F.) using the Aperio ImageScope software, with criteria outlined in (Ding and Van Hoesen, 2015) and verified

by experienced pathologist, Dr. Robert Hammond (see Appendix 3 for additional details on how these segmentations were performed). We employed our previously developed and validated pipeline for MRI and histology registration (Goubran *et al.*, 2013, 2015) to perform direct validation of 7-Tesla hippocampal subfield segmentation against ground-truth histological sections. The histology-MRI image registration procedure involved iterative 2D-3D deformable registration of downsampled (100 micron/voxel) histology slides to the reference 9.4-Tesla tissue MRI, along with 3D deformable landmark-based registration of the 7-Tesla MRI to the 9.4-Tesla MRI. Segmentation labels as well as each Laplacian gradient from the *in-vivo* 7-Tesla images were then propagated to this aligned histology space for direct comparison.

2.3 Results

2.3.1 Detection and labelling of the SRLM and hippocampal grey matter

To assess reliability, inter-rater DSI was calculated for SRLM and grey matter labels. Note that DSI tends to be lower for thin structures at higher resolutions because of the high surface area to volume ratio. DSI revealed good spatial overlap in both the SRLM (0.72 ± 0.03 right; 0.70 ± 0.04 left) and hippocampal grey matter (0.84 ± 0.01 right 0.81 ± 0.02 left). Thus, our dataset contained sufficient contrast to detect and label the SRLM and grey matter based on the visual features described in Appendix 1 with good consistency.

2.3.2 Subfield borders in unfolded coordinate space

Hippocampal subfields projected into unfolded coordinate space are shown in Figure 9. As expected, the same proximal-distal arrangement of subfields was found throughout the entire hippocampus in unfolded coordinate space, including the hippocampal head. Variability was low for all borders (i.e. low SEM compared to the area of each subfield), and no subfields crossed over each other, either in the group average or in any given unfolded segmentation example.




Figure 9. Hippocampal segmentations in unfolded coordinate space. A) Example of a manual subfield segmentation based on (Ding and Van Hoesen, 2015). Similarly, B) shows an exemplar with fewer digitations and less medial curvature. C) shows ‘unfolded’ hippocampal grey matter, with subfield label identity determined at each long-axis and proximal-distal coordinate from the manual segmentations (winner-takes-all over the sample). The shaded areas indicate standard error of the mean for each subfield boundary location across the sample of manual segmentations. Dotted lines approximately indicate commonly used boundaries between the hippocampal head, body, and tail, with the head further subdivided into uncus and vertical component of the uncus.

Reliability of the manual subfield segmentations (based on Ding and Van Hoesen, 2015) and comparison of unfolded group-average segmented hippocampi to their manually labelled counterparts are displayed in Figure 10. Lowest DSI was seen for the smallest subfields- CA2 and CA3. The data show that combining these two labels, as in some other manual segmentation protocols, leads to moderate improvements.




Figure 10. Spatial overlap in Dice Similarity Index (DSI) between manual subfield segmentations (manual-manual) and between leave-one-out unfolded group-average subfield segmentations and their manually segmented counterparts (unfolded-manual). The leave-one-out technique was performed such that borders from one participant's left and right hippocampi were not included in the averaged borders that informed unfolded segmentation of that participant's hippocampi.

2.3.3 Quantitative unfolded tissue properties

Unfolding provides a way to view grey matter properties across the entire extent of the hippocampus in a single 2D view. This unfolded view can obviate patterns that are not apparent when limited to single slices in native 3D space. Here, we mapped intracortical myelin and cortical thickness (Figure 11A and B, respectively). It should be noted that additional properties, including those used by (Glasser *et al.*, 2016), can be mapped in this way as well. We also mapped these results to the surface of a representative 3D hippocampal model in order to visualize them in native space and to allow for easier comparison with manual and unfolded average segmentations (Figure 11C).

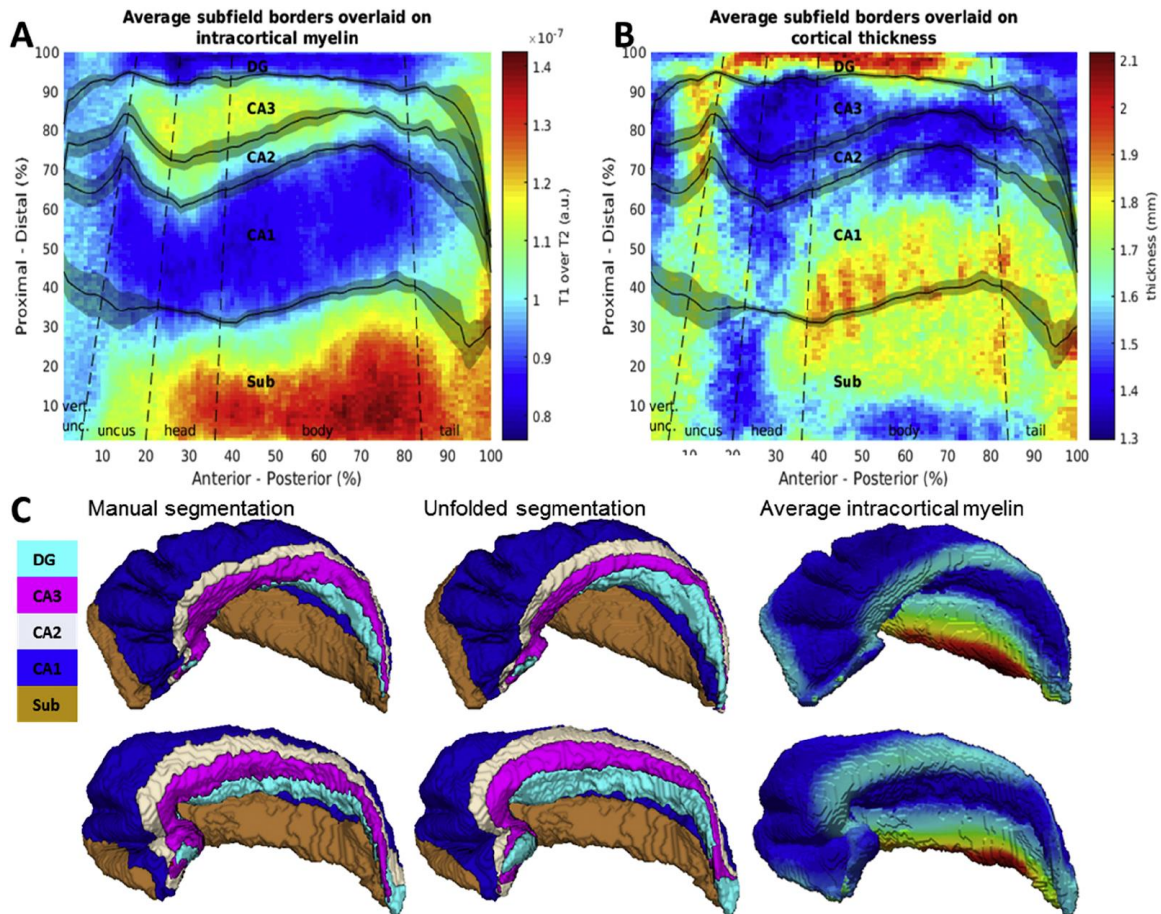


Figure 11. Quantitative mapping in unfolded coordinate space compared to subfield borders. A) Average intracortical myelin estimates (T1 over T2 MR intensities; arbitrary units). B) Average cortical thickness. Both A) and B) have average subfield borders overlaid. Note that in the dentate gyrus, thickness estimates are perpendicular to the true laminar structure (see Chapter 2.3.2). C) Manual and unfolded subfield segmentations compared to intracortical myelin in highly digitated (top) and less digitated (bottom) representative hippocampal models. Average intracortical myelin is mapped to the surface models of these hippocampi for easier comparison.

2.3.4 Histological validation

Segmentation of *in-vivo* 7-Tesla data from the surgical patient are compared to the same patient's *ex-vivo* resected and histologically stained hippocampus in Figure 12. Atrophy and cell loss in area CA1, CA3, and the dentate gyrus, with relative sparing of CA2, can

be seen in the epileptogenic tissue (Figure 12, far left). Cell loss is most evident in the distal portions of CA1, where reduced staining is seen, and atrophy is apparent in the small area occupied by each subfield compared to histological references (e.g. Ding and Van Hoesen, 2015), as well as reduced digitation structure (see Oppenheim *et al.*, 1999). These findings describe classical hippocampal sclerosis (Blümcke *et al.*, 2007). Note also that this patient shows only two clear digitations in the hippocampal head, which were not well captured in the *in-vivo* labelling of grey matter tissue (i.e. digitations cannot be seen in any of the *in-vivo* labelled data from this patient). *In-vivo* segmentations showed some misalignment of both grey matter and SRLM labels, which can be seen in areas where neurons are visible in the histology without being obscured by grey matter labels. This likely reflects imperfect alignment between the *in-vivo* scan and *ex-vivo* histology but may also be due to poor image quality in this patient, making it difficult to correctly label hippocampal grey matter and SRLM.




Figure 12. Neu-N stain of resected hippocampal tissue with comparison of histologically segmented hippocampal subfields to *in-vivo* subfield labels and unfolding Laplace gradients in the same individual. Slices move from posterior (hippocampal body) to anterior through the hippocampal head and are separated

by 4.4mm. Red arrows indicate marked cell loss in distal CA1. The proximal-distal and long-axis gradients are surrounded by low colour-scaled voxels because of their interpolation when transforming to the histological space.

Some inconsistencies are noticeable in direct comparison between the *in-vivo* manual subfield segmentation (Figure 12 column 2) and the histological ground truth (Figure 12 column 3). They may be due to misalignment, tissue atrophy, poorer image quality, errors in the segmentation protocol, or inter-individual differences in subfield border locations. Note, however, that some key features are preserved in our segmentations: CA2 passes into plane twice in row 3 (i.e. appears in two different places), and both dentate gyrus and CA3 are seen in the uncus in row 2. Because these borders are curved medially in the head, they are difficult to capture in coronal slices alone and will vary drastically depending on the exact position of the slice.

The long-axis and proximal-distal gradient together make up our two-dimensional unfolded coordinate space, which can be viewed in relationship to histology in Figure 12 columns 4 and 5. The long-axis gradient changes between coronal slices, as expected when moving from posterior to anterior, but also changes within each slice. This is because the anterior point is located in the most medial, vertical component of the uncus, where each of the subfields has its natural terminus, rather than the absolute anterior of the hippocampus. Thus, the long-axis gradient shows how any given coronal slice is out-of-plane with respect to the medial curvature of the uncus. The proximal-distal gradient identifies a set of potential subfield borders, which can also be adjusted depending on anterior-posterior extent. In row 3 this gradient can be seen to increase and then decrease as the gradient passes into and then out of the plane of view (i.e. from proximal to distal colours pass from green-yellow-orange-yellow-orange), similar to what is seen for CA1 and CA2 in the corresponding histological ground truth images. We used this gradient in combination with the long-axis gradient to apply group-averaged subfield borders to this participant's unfolded hippocampal space, which is shown in the rightmost column. This segmentation suffers from some of the same issues as the fully manual segmentation, but as in the manual segmentation, many of the key features of the hippocampal head are retained, including the passing into and out-of-plane for CA2 and the presence of CA3 in

the uncus (Figure 12 rows 2 and 3, respectively). However, in the unfolded group-average segmentation and in the fully manual segmentation, each of the *in-vivo* segmentation borders was incorrectly placed. Specifically, all subfield borders should be shifted more distally in rows 1 and 2, and in rows 3 and 4 the manually segmented border shifts are mixed whereas the group-averaged borders should be shifted more proximally. This resulted in no CA1 label for the slice seen in row 4 of the unfolded group-average segmentation, despite its presence in the histological ground truth.

2.4 Discussion

Using isotropic T2 7-Tesla MR imaging, we were able to detect substructures of the hippocampus that can be leveraged to understand and quantify its complex and variable topology. Towards this end, we developed a methodology to ‘unfold’ or index this topology in a way that respects accounts of hippocampal subfields from the literature and inherently aligns tissues despite variable folding. The resulting system can be used to index or segment that tissue with a high level of precision and flexibility. Segmentations performed in the resulting unfolded space agree with quantitative metrics of intracortical myelin and capture key features shown in an *ex-vivo* validation sample from a patient with epilepsy.

2.4.1 Detection and labelling of the SRLM and hippocampal grey matter

High Dice Similarity Indices were found for both hippocampal grey matter and SRLM using the segmentation instructions outlined in Appendix 1. This feature is critical in that it allows for the differentiation of folds throughout the entire hippocampus and was necessary to constrain the proximal-distal Laplacian solution obtained here to the topology of hippocampal archicortex. Another constraint worth mentioning is that, in our protocol, the SRLM is required to separate folds of hippocampal grey matter with no points of contact between the folds. This feature may be challenging to obtain in datasets generated with highly anisotropic acquisition, which will limit visibility of the SRLM in the medial extensions of the hippocampus, and also limit the rater's ability to create a label with separated folds of hippocampal grey matter (as they may change too drastically

between slices). However, with sufficiently thin slicing, it may still be possible to employ these criteria and procedures in anisotropic datasets.

2.4.2 Anatomical details of unfolding

The following list describes specific anatomical details that were correctly captured with our unfolded coordinate system:

- In the long-axis unfolding, a cross-section of hippocampal grey matter at an equipotential point in the hippocampal body reveals the same classic C-shaped orientation of grey matter as in coronal slices from histology and in extant MRI tracing protocols (Figure 8A, first right panel, red dotted line). Unlike in extant protocols based on coronal or other views, cross-sections in our unfolding at equipotential points in the hippocampal head also correctly reveal these C-shaped orientations of subfields (e.g. Figure 8A, second right panel, red dotted line).
- In the vertical component of the uncus, the C-shape of hippocampal grey matter flattens out to form a line (e.g. Figure 8A, third and fourth right panel, red dotted line). Here, the dentate gyrus passes most medially around the other subfields before extending upwards and reaching the vertical component of the uncus (Ding and Van Hoesen, 2015). This feature is accounted for in our unfolding by the manual placement of the ‘sink’ in the proximal-distal unfolding of hippocampal grey matter (Figure 8B, third and fourth right panels, turquoise bin; see also Appendix 1, step 3: labelling of extra-hippocampal structures).
- In the proximal-distal unfolding, the more proximal regions of hippocampal grey matter wrap around the absolute anterior tip of the hippocampus, moving from the inferior to superior side. This feature honours the descriptions of the subiculum provided by Ding and Van Hoesen (2015) and discussed in Chapter 2.1.2 (Figure 8B, main model, green, blue, and yellow and cyan bins following dotted red arrow)

There are some remaining anatomical limitations of the current unfolded coordinate system:

- The ‘sink’ used for proximal-distal unfolding captures most of what corresponds to the granule cell layer of the dentate gyrus in histological studies. However, this tissue is so thin that it cannot easily be matched to the true granule cell layer seen in histology. Thus, we do not recommend using it as an independent region of interest for *in-vivo* MRI. Rather, we recommend combining it with area CA4, or the CA3 hilar region, as implemented in our manual subfield segmentations and in our unfolded subfield descriptions.
- The laminae of the dentate gyrus are known to be situated perpendicular to those of the other subfields (Duvernoy *et al.*, 2013). This is not respected by our laminar unfolding. Instead, most of the dentate gyrus is treated as being deep laminae (e.g. Figure 8C, second right panel, correct lamination shown in white dotted line). Thus, caution is necessary when the goal is to index the laminae of the CA4 and dentate gyrus.

Another strength of this ‘unfolded’ space pertains to the fact that all its distances are relative to the full size of the corresponding individual hippocampus and can thus be applied across a range of hippocampal sizes and morphologies. With some adaptation of landmarks used as boundaries in the Laplace equation (i.e. source and sink), this protocol may also be applied to the characterization of abnormal hippocampi (due to abnormal development or neurological disease) or even those from other mammalian species.

2.4.3 Subfield borders in unfolded coordinate space

Figure 9 shows the mapping of subfields segmented in each participant's native space to the standardized unfolded space. The fact that the SEM of the average unfolded borders was relatively low (i.e. accounts for a relatively small proportion of the area of each subfield) is surprising, given the large inter-individual variability of subfield locations in native space (e.g. highly variable digitations; see Ding and Van Hoesen, 2015). This result suggests that much of the variability in native space is due to differential curvature and folding of hippocampal tissue in development, rather than differences in the cytoarchitectural differentiation within this tissue.

Comparison of manual segmentations in native space to segmentations applied using the group-averaged borders in unfolded space revealed moderate to good spatial overlap, as determined by DSI, particularly when the smallest subfields CA2 and CA3 were combined (Figure 10). These DSI scores were also similar to inter-rater manual DSI scores. However, the sources of the remaining variability are not clear. They might reflect individual variability in subfield border locations that are not captured by our unfolded average borders. Alternatively, they might reflect deviations from the true subfield border locations in a manual segmentation in native space due to tracing errors.

2.4.4 Quantitative unfolded tissue properties

An important finding for the intracortical myelin estimates we obtained is that they appear to closely correspond to the average subfield borders used in our demonstration of subfield segmentations (Figure 11A). The subiculum and areas CA2 and CA3 appear to have greater myelin content than CA1 and the dentate gyrus, with the proximal part of the subicular complex showing greatest values. Though speculative, we suggest that this characteristic could reflect contributions of the perforant path passing through the subiculum, elevating myelin estimates due to the presence of white matter tracts. In area CA3, dense recurrent collaterals might contribute to elevated myelin estimates. An alternative explanation is that increased vasculature, which would also appear dark in T2-weighted images, contributes to this contrast. Support for this explanation comes from the observation that area CA2 is the most highly vascularized subfield in humans (Duvernoy *et al.*, 2013). Our findings also agree with those of Ábrahám *et al.* (2012) who examined intracortical myelin in histological samples, and those of Marques and Gruetter (2013), who found similar differences in R1 MR intensities between the subfields.

Cortical thickness was also calculated using Laplacian streamlines and mapped in unfolded coordinate space (Figure 11B). Critically, these differences do not appear to correspond to the subfield borders. Note that thickness in the dentate gyrus was actually calculated perpendicular to the true laminae of the dentate gyrus because of the different orientation of tissue (see Chapter 2.3.2). Therefore, thickness in this region should not be considered to reflect true laminar structure. With this exception, our results are similar to thickness measures reported by Yushkevich *et al.* (2011), who found that thickness was

highest between subiculum and CA1 and lowest in CA3, both in healthy individuals and those with mild cognitive impairment. However, these results differ from those reported by Burggren *et al.* (2008), who found that thickness was highest in CA3/dentate gyrus, particularly in the anterior hippocampus. This discrepancy may be due to the fact that their method did not account for the digitations in the hippocampal head, which could have led to overestimation of thicknesses.

Comparing both manual and unfolded group-average subfield segmentations performed in this study to intracortical myelin, we noted some discrepancies between border locations. In the representative hippocampi examined, the subfield borders in the manual segmentation did not follow as smooth of a trajectory as in the unfolded segmentation (Figure 11C). This discrepancy could be due to unique aspects of this particular participant's subfield border locations that may not be captured by the group average. However, we believe it is more likely due to limitations of the manual segmentation employed. For example, imperfect alignment between coronal slices in MRI with histological reference slices from Ding and Van Hoesen (2015) and Duvernoy *et al.* (2013) could cause subfield borders to shift, making them more jagged when in reality they follow a smooth trajectory. Detailed 3D histological examinations are needed to determine whether this is indeed the case. These issues may contribute to the variability observed between manual segmentations and unfolded group-averaged segmentations that are apparent in Figure 10. They also speak to more fundamental challenges related to the reliability and feasibility of manual segmentation protocols that are based primarily on coronal slices, which fall beyond the scope of the current paper. Nevertheless, the fact that these discrepancies largely average out in a sample of unfolded hippocampi, while still respecting differences in morphology, further highlight the strength of the approach presented here.

The unfolding of hippocampal grey matter accounts for much of the inter-individual variability related to differences in ontological folding. However, inter-individual differences in subfield border locations beyond differences in folding structure (for example, as shown by Zeineh *et al.*, 2015), as well as variability due to the presence of disease (as in the resected tissue examined here) may reflect additional sources of

variability in unfolded subfield locations and size. This variability presents a significant challenge when aiming to apply unfolded group-average borders, as well as for manual or automated segmentation protocols that rely on geometric rules and structural landmarks. Thus, the use of other cues, such as intracortical myelin or thickness, may be useful in generating subject-specific subfield borders in future follow-up research.

2.4.5 Histological validation

Results from Figure 12 show some inconsistencies in our manual subfield segmentation as well as in our group-averaged unfolded space segmentations when compared to the histological ground truth in one resected tissue sample from a patient with epilepsy. In addition to general segmentation errors and natural inter-individual variability in border locations, some of these inconsistencies may arise because of tissue atrophy and poor grey matter labelling in this participant with a neurological condition as compared to healthy control participants. Such factors may make manual segmentation based on Ding and Van Hoesen (2015)'s descriptions, or the use of borders established in healthy control participants less appropriate for characterization of the hippocampi in a disease state. However, it should be noted that several features seen in the histology were still captured by our unfolding coordinate system, despite being absent from the subfield borders applied here. In particular, the proximal-distal gradient can be seen to increase and decrease along the length of a coronal slice, capturing how the CA1 passes into plane twice in the histological ground truth. Thus, although subfield borders may be shifted due to various factors, the coordinate space presented here still respects this feature of the hippocampal head. As such, this application illustrates how segmentations in unfolded coordinate space are able to capture critical structural complexities of the hippocampal head.

2.4.6 Hippocampal unfolding in the context of extant literature

One possible source of the recent controversy over hippocampal subfield borders relates to the constraint that in order to be reliable, a coronal slice segmentation protocol should make use of heuristics, such as geometric rules with reference to visible intra- or extra-hippocampal landmarks. For these rules to be applicable across different hippocampal

morphologies and MR image qualities, some level of simplification is necessary, reducing accuracy and precision. Given the large number of possible border locations, the unfolded coordinate system presented here inherently allows for increased precision, even across hippocampi with varied morphologies. This is because it respects critical structural features of the hippocampal head without reliance on the heuristics mentioned. Thus, although we do not wish to present the specific subfield borders used here as an alternative to the efforts towards international harmonization by the Hippocampal Subfields Group, we hope that these efforts will include the structural considerations discussed in the current paper, and may also lead to exploration of methodologies other than manual segmentation. Furthermore, we anticipate that, once international consensus is reached, the resulting subfield borders can be applied using the unfolded coordinate system presented here and be complemented by further characterization of inter-individual differences that can be captured with the present methodology. Given the increasing prevalence of high-resolution data in which the SRLM can be identified throughout the length of the hippocampus, this appears to be a particularly promising avenue.

A final point worth noting is that the unfolded coordinate system offered here will also allow for easy implementation of further subfield divisions in future work. For example, Ding & Van Hoesen's recent characterization of the hippocampal head (2015), as well as other histological evidence from humans and nonhuman animals (see Ding, 2013), reveal differentiation of the subiculum into distinct components, including the prosubiculum (postsubiculum in rodents), subiculum, presubiculum, and parasubiculum. Furthermore, some studies have documented functional differentiation between proximal and distal CA1 (Nakazawa *et al.*, 2016, Knierim *et al.*, 2014) and CA3 (Nakamura *et al.*, 2013). These findings highlight the increasing need for precision and standardization in indexing hippocampal tissue, as well as the need for flexibility in applying subfield labels so as to honour new developments in tissue characterization. We believe that the unfolded coordinate system presented here can provide such a framework.

2.4.7 Conclusions

We have presented a new tool that promises to allow for *in-vivo* characterization of the complex structure of the human hippocampal subfields in unprecedented detail. Manual segmentation with high anatomical detail poses many challenges for the generation of reliable protocols that are suitable for tracing of hippocampal subfields, particularly in the hippocampal head. However, consideration of regularities in hippocampal structure related to ontogeny offer ways in which computational tools, such as the Laplace equation, can be applied for indexing and segmenting hippocampal tissue in a way that preserves topology across individual differences. In the current study, we pursued an approach that took advantage of these considerations. Through computational unfolding of the hippocampus, the current protocol provides a coordinate system that can index hippocampal tissue in a precise and flexible manner, while capturing the noticeable inter-individual differences in morphology that have been documented in histological studies of this structure. This method critically depends on the visualization of the SRLM, or ‘crease’, along which the hippocampus is folded. We argue that this method offers several practical advantages over manual segmentation techniques. These advantages can be summarized as follows:

- Unfolding hippocampal grey matter allows for indexing of analogous tissues (or sets of candidate boundary locations) across participants with variable morphologies.
- The unfolded coordinate space can be used for inter-subject alignment and subsequent mapping of properties across the full long-axis and proximal-distal extent of the hippocampus, as illustrated here for intracortical myelin and cortical thickness measures.
- Segmentations applied in this unfolded coordinate space show good spatial overlap with, and may even correct for tracing errors in, detailed manual subfield segmentations. This coordinate system also captures subtle but critical structural features, as demonstrated in a direct comparison with a resected histological sample from a patient with epilepsy.

Future directions for this work include the integration of automatic tissue segmentation tools for detection of the SRLM, grey matter, and surrounding structures in order to reduce user input and improve reliability. Promising applications of this unfolded coordinate system include cross-species comparison and normative mapping of hippocampal tissue properties in health and disease.

Chapter 3

3 Unsupervised 3D Histology Analysis

This Chapter focuses on a data-driven analysis of the open source BigBrain 3D histology dataset under the hippocampal unfolding framework presented in Chapter 1. This work is now published in the journal *NeuroImage* under the title ‘Hippocampal subfields revealed through unfolding and unsupervised clustering of laminar and morphological features in 3D BigBrain’ (DeKraker *et al.*, 2020). BigBrain consists of an *ex-vivo* brain that was serially sectioned, stained, imaged, and then reconstructed in 3D at $40\mu\text{m}^3$ resolution (Amunts *et al.*, 2013). We performed manual tissue segmentation of the left and right hippocampi with unprecedented detail, then applied our previously developed unfolded coordinate framework with minor changes to the code. A set of 10 laminar features and 5 morphological features were computationally extracted in unfolded space, and data-driven clustering was performed on these features. Resulting clusters closely resembled manual subfield definitions from the native 3D histology. This speaks to the sensitivity of the feature extraction methods applied here, as well as the validity of subfield definitions used. Translation, insights, and limitations of *in-vivo* neuroimaging are discussed as they pertain to the observation made in this project.

3.1 Introduction

The hippocampus is one of the most heavily investigated brain structures in neuroscience. Much research in recent years has focused on questions about its subdivisions, guided by the idea that different regions within the hippocampus may perform different functions and may also be differentially prone to disease (Small *et al.*, 2011). These developments pose central questions as to how to characterize subdivisions in anatomical terms.

Traditionally, most proposed subdivisions have relied on histology and cytoarchitecture, leading to the notion of distinct hippocampal subfields that typically include the subicular complex, Cornu Ammonis 1 to 4, and the dentate gyrus (Duvernoy *et al.*, 2013). More recently, increasing interest has also emerged concerning graded differences along the anterior-posterior axis based on subfield composition and connectivity (Strange *et al.*, 2014; Poppenk *et al.*, 2013; Plachti *et al.*, 2019). An organizational principle that shapes

these dimensions, i.e., subfields and anterior-posterior differences, is the complex topology within the hippocampus that results from folding during its ontological development (Duvernoy *et al.*, 2013; DeKraker *et al.*, 2018). This principle has received only limited investigation to date but requires careful consideration in any effort to characterize the internal architecture of the hippocampus. The current paper aims to investigate the relationship between hippocampal morphology and laminar cytoarchitecture under a topological framework in humans. In other words, the goal is to examine laminar and morphological (i.e. non-laminar) features, such as cortical thickness or curvature, within a framework that explicitly honours 3D continuity across 2D images. In order to pursue this goal, we took advantage of the unique and powerful “BigBrain” dataset that provides continuous histological sampling with full 3D coverage (Amunts *et al.*, 2013). A particular promise of this approach lies in its applicability to *in-vivo* Magnetic Resonance Imaging (MRI).

While commonly used MRI measures do not allow for cytoarchitectural characterization, MR-based protocols have been developed to indirectly infer the locations of hippocampal subfields in humans based either on manually delineated landmarks or corresponding probabilistic atlases that are informed by histological reference material (Yushkevich *et al.*, 2015a, 2015b; Iglesias *et al.*, 2015). However, traditional histological references can be problematic for several reasons. First, they often contain only select coronal slices taken from regions where folding is the simplest, most frequently from the hippocampal body, with the notable exception of (Ding and Van Hoesen, 2015) who focus greater attention on the hippocampal head than most other investigations. Second, even in the hippocampal body slices are taken sparsely, limiting the number of contextual features that can be gathered from neighbouring slices or other planes of view. Third, histological preparation often deforms the tissue of interest relative to its *in-vivo* state, which is a problem for MRI co-registration unless the histological sample is also imaged prior to histological preparation. Finally, even among neuroanatomists there is some disagreement as to exactly which labels, stains, and histological features should be used for defining hippocampal subfields (Wisse *et al.*, 2017). Some previous studies have made use of *ex-vivo* MRI to aid in the translation of histology to MRI (Iglesias *et al.*, 2015; Yushkevich *et al.*, 2009) in an effort to mitigate some of these issues. However,

even with such an approach, inter-individual differences in hippocampal morphology can impose limitations for inferring subfields or other structural features, when hippocampal topology is not considered.

It is well established that the human hippocampus is a folded component of archicortex that is continuous with the neocortex (Duvernoy *et al.*, 2013; Nieuwenhuys *et al.*, 2013). The hippocampal folds include wrapping around its innermost region - the dentate gyrus, as well as anterior-posterior folding that is sometimes referred to as dentation, digitation, or gyrification. The gyrification seen in the hippocampus is morphologically similar to gyrification in the neocortex (although not necessarily based on the same ontogeny). It has been shown to vary considerably between individuals (DeKraker *et al.*, 2018; Chang *et al.*, 2018) and can be affected by age (Cai *et al.*, 2019) or disease, such as temporal-lobe epilepsy (Blümcke *et al.*, 2013). This folding is an important aspect of understanding the internal structure of the hippocampus, and for appreciation of the continuity of subfields, particularly in its anterior portion that includes the uncus (Ding and Van Hoesen, 2015). Topological analyses can provide a framework for extracting these continuities, for example through unfolding (DeKraker *et al.*, 2018), and offer the basis for laminar and further morphological characterization of complete hippocampal structure in 3D, including subject-specific gyrification.

The dataset made publicly available by BigBrain (Amunts *et al.*, 2013) provides a unique opportunity to conduct topological analyses of histology data in 3D, and to examine topological measures in unfolded tissue. This dataset consists of 3D histology, digitally reconstructed from images of serially sectioned and stained cadaveric brain tissue. In the current project, we used reconstructed blocks of the left and right hippocampi (40 μ m isotropic) from BigBrain to identify topologically-derived laminar and morphological features under our hippocampal unfolding framework. To characterize laminae, we focused on 10 computationally derived features describing the distributions of neurons (Amunts *et al.*, 1999), which were also recently used to characterize the neocortex in BigBrain (Wagstyl *et al.*, 2018; Wagstyl *et al.*, bioRxiv). Morphological features were also computationally derived and included thickness, curvature, inner and outer surface textures, as well as gyrification. We then compared these morphological and laminar

features to classic descriptions of subfields and examined variations along the anterior-posterior hippocampal axis. We anticipated that the features examined would differ substantially between subfields. Therefore, we also tested whether it might even be possible to obtain successful subfield segmentation with an unsupervised feature-based approach. This type of approach is desirable for its objectivity, which could help resolve differences among neuroimagers and histologists on subfields definitions. It also allowed us to examine which subsets of features are sufficient to derive clusters resembling gold standard hippocampal subfields. For this purpose, we contrasted the contributions of laminar and morphological features, given that laminar features are used most prominently in histology (see Duvernoy *et al.*, 2013; Nieuwenhuys *et al.*, 2013) but morphological features, such as thickness, are more readily available in high-resolution structural MRI (e.g. DeKraker *et al.*, 2018).

3.2 Methods

The backbone of our analyses was to impose a topological unfolding framework to manual hippocampal traces, a method that we previously developed for 7-Tesla MRI (DeKraker *et al.*, 2018). We then extracted various morphological features of hippocampal structure from the left and right BigBrain hippocampi. We computed laminar features based on the work of (Amunts *et al.*, 1999) and modelled as in (Wahnert *et al.*, 2014). We then performed unsupervised, data-driven clustering of these features and compared resulting clusters to manually segmented hippocampal subfields. Finally, we examined differences in hippocampal structure along its longitudinal (i.e., anterior-posterior) axis.

3.2.1 Materials

Histological data used in this study came from the BigBrain dataset, which consists of serially sectioned and stained brain tissue that was then reconstructed in 3D. Specifically, in the present study we used bilateral $40\ \mu\text{m}^3$ resolution hippocampal blocks (ftp://bigbrain.loris.ca/BigBrainRelease.2015/3D_ROIs/Hippocampus/) in addition to serial section images at $20\ \mu\text{m}^2$ resolution (ftp://bigbrain.loris.ca/BigBrainRelease.2015/2D_Final_Sections/Coronal/Png/Full_Reso

lution/) (Amunts *et al.*, 2013). BigBrain preparation involved silver staining, which selectively targets cell bodies, providing a contrast that is different from, but conceptually similar to, Grey Level Index (Schleicher *et al.*, 1999; Wagstyl *et al.*, 2018). Because of the large file sizes, tracing and application of our unfolding framework were performed on downsampled images (80 μ m isotropic) before upsampling by nearest-neighbour interpolation in the case of labelmaps, and linear interpolation in the case of unfolding solutions.

3.2.2 Manual tracing

Detailed histological tracing was performed for each hippocampus by a combination of manual tracing and the user-guided computational tools in ITK-SNAP 3.6 (Yushkevich *et al.*, 2006). ITK-SNAP is primarily used for manual tracing using a 2D or 3D paintbrush, but it additionally contains some semi-automated tools that were used at the manual tracer's discretion, including morphological operations such as dilation and erosion. Since these tools rely so closely on the supervision of the tracer, we will refer to the use of these tools as a manual process throughout this manuscript. All traces were performed in native space, using all three view planes. A general label for hippocampal grey matter (subiculum and CA1-4) was manually traced first, and this tissue was later manually divided into subfields. Only the laminae which contained stained neuronal cell bodies - stratum pyramidale, oriens, and lucidum - were traced (Figure 13). Strata radiatum, lacunosum, and moleculare (SRLM) and the alveus were given separate labels even though they are sometimes considered laminae of the archicortex containing dendrites and axons of pyramidal cells (Duvernoy, Cattin and Risold, 2013); they were not stained by this contrast (although note that some of these strata contain interneurons - see Appendix 4 for discussion).




Figure 13. Manual traces of hippocampal archicortex and segmentation into subfields. A) shows coronal slices through the left hippocampal head (rows 1–3) body (row 4) and tail (row 5), with manual segmentations overlaid in the images to the right. Images were taken from coronal slices 716, 632, 590, 376, and 230 of the 40 μm native space left hippocampal block. B) shows 3D models of each hippocampus as seen from their superior aspect, with the inferior aspect shown in C). Dotted lines in B) indicate approximate locations of each coronal slice shown in A). SRLM, vestigial hippocampal sulcus, alveus, and fimbria were excluded from all labels. Red arrows indicate anterior folding in the vertical component of the uncus, the orange inset in the third row highlights ‘islands’ of neuronal cell bodies in the subicular stratum lacunosum, and yellow arrows indicate gyrifications in the posterior body and tail of the hippocampus.

Subfield segmentation (i.e. the division of archicortical grey matter into distinct subfields) was performed in 3D by rater Kayla Ferko (K.F.) according to the criteria outlined by (Ding and Van Hoesen, 2015). This work provides a detailed and instructive guide to segmenting all subfields of the hippocampus, including the hippocampal head with multiple samples with varied folding structure. One limitation of our segmentations is that Ding *et al.*'s protocol includes the use of both neuronal body and myelin stains, whereas no myelin stain is available in the BigBrain dataset. To take full advantage of the histological features available in BigBrain and make use of the highest resolution available, original 20 μm images were also consulted every 2 mm. In other words, subfield borders on the 80 μm 3D hippocampus were manually compared to borders on 20 μm 2D images and, where necessary, the 80 μm borders were updated to best match 20 μm images. These segmentations included the subiculum and CA1-4 but did not differentiate the regions within the subicular complex due to lack of resolution and since no myelin staining was available. Subfields were traced through the entire length of the hippocampus including the uncus and vertical component of the uncus, in which (Ding and Van Hoesen, 2015) describe modified versions of the same subfields. Because the vertical component of the uncus is very thin, its subfields were not easily discriminable. As a consequence, they were partially inferred from neighbouring regions of the hippocampus. Smoothing of the final traces was then performed by label-specific dilation (e.g. dilating subiculum 4 voxels over CA1; then dilating CA1 8 voxels over subiculum; then dilating subiculum 4 voxels over CA1 again, while always ensuring dilation was balanced in both directions). This was performed for all subfields in the order CA4-CA3-CA2-CA1-subiculum, ensuring that subfield borders followed smooth transitions from slice to slice.

Structures surrounding the hippocampus were traced only in the regions that border the hippocampus. These labels included medial-temporal lobe neocortex (MTLc) (entorhinal and parahippocampal regions), hippocampal-amygdalar transition area (HATA), and indusium griseum (ind. gris.). HATA borders were clearly discriminable from archicortex by a marked change in density and physical separation from archicortical neurons. Ind. gris. and MTLc borders were less clear, and so they were demarcated using the heuristics used in previous work in MRI (for example, the MTLc-subiculum boundary is defined at

the super-medial edge of the white matter in the parahippocampal gyrus. See (DeKraker *et al.*, 2018) or Chapter 2 for further details).

3.2.3 Topological unfolding framework

In previous work (DeKraker *et al.*, 2018 or Chapter 2), we imposed a topological unfolding framework on the hippocampus by solving Laplace's equation over the domain of the hippocampus under multiple sets of boundary conditions: anterior-posterior, proximal-distal, and laminar. The anterior-posterior and proximal-distal solutions can then be used to index regions of the hippocampus in 2D according to its topology, irrespective of inter-individual differences in gyrifications, rotation, curvature, size, orientation, or position of the hippocampus. This provides implicit registration between hippocampi despite inherently different morphologies. Data can be transformed between native and unfolded space bidirectionally using interpolation. We applied this same approach to BigBrain hippocampal traces (see Figure 14 for illustration). However, note that several minor improvements were made to this code which are detailed in Appendix 5. Most notably, instead of binning maps into 100×100 pixels, they were instead derived from a mid-surface mesh consisting of 512×256 vertices, with data being sampled from all features by nearest neighbour. The dentate gyrus (DG) was not included in this unfolding. Although it was easily distinguishable from other subfields by its very high cell density, it is topologically disconnected from the rest of the archicortex, and therefore would be out-of-plane (i.e. perpendicular) to our unfolded space (see Figure 13 for visualization). Note that this differs from the work presented in Chapter 2, which presents the thickness of the DG in what would conventionally be considered the width of the DG granule cell layer.




Figure 14. Topological unfolding framework in BigBrain with hippocampal subfields. A) Sagittal slice and 3D models of the Laplacian solutions (proximal-distal and anterior-posterior) for the right hippocampus. Image was taken from sagittal slice 514 of the 40 μm native space left hippocampal block. B) Mid-surface topological models of the left and right hippocampi in native and unfolded space.

Waehnert *et al.* (2014) noted that neocortical laminae are displaced due to curvature in gyri and sulci, and they propose an ‘Equivolume’ model that captures this feature better than a Laplacian (or equipotential) solution. Their model is motivated by the observation that a given lamina, for example near the pial surface, will stretch at the apex of a gyrus and compress at the depth of a sulcus, causing it to become thinner and thicker in these respective regions, and vice versa for laminae at the white matter surface. Thus, we also included an alternative laminar indexing system using the Equivolume model solution obtained from Nighres (Landman *et al.*, 2013). Again, this was performed on the downsampled (80 μm) traces before upsampling as described above. The resulting model had fewer gyrification-related artifacts in laminar profiles and was used for all subsequent laminar analyses. However, some other artifacts were observed under this model solution, likely as a result of the rough texture of the subiculum surface (see Appendix 6 for details). These laminar profiles were extracted for each unfolded point (512 \times 256 \times 16 points, or vertices) at the corresponding nearest neighbour (full resolution) native space voxel.

3.2.4 Morphological feature extraction

Each morphological feature is illustrated in the top left panel of Figure 15. Thickness estimates were obtained across the unfolded space of the hippocampus as in previous work, that is, by generating and measuring streamlines in 3D across the laminar Laplacian solution obtained from our topological unfolding framework. Curvature estimates were obtained by generating a mid-surface along the hippocampus with the vertices being interpolated xyz coordinates from each unfolded point at a laminar distance of 0.5, which is the midpoint between the inner and outer surface. Smoothing of face normals was applied, and mean curvature was then estimated at each vertex (see Appendix 5 for details). The inner (i.e. adjacent to the SRLM; continuous with the neocortical pial surface) and outer (i.e. adjacent to the alveus, continuous with the neocortical white matter surface) surfaces of the hippocampus were rougher than their mid-surface counterpart due to the presence of other features, such as subicular ‘islands’ of cell bodies shown in Figure 13. Thus, we additionally computed curvatures of these surfaces after smoothing as described above. Gyrification is typically defined as a ratio outer surface area, for example that of a brain mask over gyrified surface area, in this example including sulcal area (Larsen, Nielsen and Sparring, 2006). Since the hippocampus is an open-ended cortical surface it does not map easily to an outer surface area or to a sphere as in the neocortex, and so our unfolding framework instead maps it to a rectangle. We thus defined gyrification as a ratio of native space surface area over unfolded surface area at each unfolded point.




Figure 15. Characterization of the hippocampus using morphological and laminar features. The top diagrams illustrate how each feature is derived (see Chapters 3.2.4 and 3.2.5 for details). Top left shows an example segment of cortex, while the top center and top right show an example laminar profile and its absolute derivative (Abs.Deriv), respectively (Amunts *et al.*, 1999; Larsen, Nielsen and Sparring, 2006). Heat maps below show the z-scored values of each feature across the unfolded hippocampus in the left and right hemispheres, with the same colour scaling in both hemispheres. Overlaid in white are manually defined subfield borders, the top edge being the border with the DG which is out-of-plane.

3.2.5 Laminar feature extraction

We extracted laminar profiles along the Equivolume laminar solution described above, and then summarized these profiles using the same 10 features consistently used by (Amunts *et al.*, 1999; Wagstyl *et al.*, 2018). Briefly, this involved sampling staining intensities (y) along a laminar profile through the cortex and calculating the mean ($\text{Mean}(y)$). This intensity profile was then treated as a distribution (x), and the mean ($\text{Mean}(x)$) and first 3 moments ($\text{SD}(x)$, $\text{Skew}(x)$, and $\text{Kurt}(x)$) were calculated. The absolute value of the derivative (Abs.Deriv) of the profile was then calculated ($y \rightarrow y.d$), and the same measures (e.g. $\text{Mean}(y.d)$, $\text{Mean}(x.d)$, etc.) were obtained. These methods are illustrated with corresponding terminology at the top of Figure 15.

There were several methods developed for 3D MRI, which we were able to incorporate into this analysis, with resulting differences when compared with the analyses performed by (Amunts *et al.*, 1999). Firstly, we sampled laminar profiles under the 3D Equivolume model that minimizes distortions in laminae due to curvature (as discussed above). Secondly, our laminar sampling was not as dense because of the reduced resolution available in the current data, and the fact that the laminae of the archicortex are generally thinner than those of the neocortex. Lastly, we included only laminae containing neuronal cell bodies (as discussed above). Further details on these differences between our methods and those of (Amunts *et al.*, 1999) can be found in Appendix 5.

3.2.6 Unsupervised clustering

In order to cluster visually homologous regions of the feature maps into segments, we applied a scale-space representation employing image pyramids. That is, for each of the selected features, we smoothed the data in unfolded space with a Gaussian kernel and a Laplacian of Gaussian kernel of sizes $\sigma = 0.16, 0.32, 0.64, 1.28, \text{ and } 2.56$ mm in order to capture features at various spatial scales. Because unfolded space does not necessarily have correspondence to real-world size, we reparametrized our unfolded space according to real-world distances between points prior to smoothing, and then returned the resulting smoothed feature maps to the original unfolded space parameterization. See Appendix 7 for details and visualization of this reparameterization.

This multi-scale smoothing is meant to reduce noise in the data, but also to provide some spatial regularization of features, ensuring that resulting clusters will not be distributed in only small patches. The anterior 10% and posterior 10% of each feature were discarded due to high noise.

All morphological and laminar features from the left and right hippocampi were then reshaped into single vectors, z-scored, and entered into a principal component analysis (PCA). K-means clustering was then computed on the first 8 components, which explained >1% variance each, with a fixed number of output clusters of $k = 5$ (since manual segmentations contained 5 subfield labels). PCA followed by K-means clustering was ideal for this type of analysis for several reasons: 1) Co-linearity among features can be clearly assessed using PCA prior to k-means clustering; 2) Clusters were expected to be of comparable sizes, which k-means is biased towards; and 3) The number of clusters is known *a-priori*. Clusters were then assigned subfield labels based on highest overlap. Dice Similarity Index (DSI) was calculated (Dice, 1945; Sørensen, 1948) in unfolded space for each subfield (i.e. disregarding thickness), excluding the 10% anterior and posterior edges that were removed due to high noise. DSI was also calculated in native space, which is shown in Appendix 7. In this case, however, clusters had to be extrapolated over the 10% anterior and posterior regions leading to lower total overlap scores. We also explored clustering under $k = [2,4,8,16,32]$, and performed alternative clustering methods, in order to determine the consistency of subfields and sensitivity to further subdivisions in the data, with results described in Appendix 7. In brief, these results support the validity of the clustering methods we employed, and show that when other methods are used, cluster boundaries occur at locations similar to those reported in the main body of this article. PCA variance explained per component, component loadings and visualization of the first 8 components can be viewed in Figure 17, along with the correlations between all features.

In order to determine whether subfield clustering could be derived using only laminar features or only morphological features alone, we repeated the above process for the subsets of morphological and laminar features separately. We used the same '>1% variance explained' threshold to remove PCA 'noise' components, which resulted in 5

components in the laminar feature clustering, and 3 components in the morphological feature clustering. Morphological features of inner and outer surface textures were excluded since they capture subicular 'islands' of cell bodies in stratum lacunosum, which could be considered a laminar feature. This exclusion was also based on the limited value of these two features for any MRI assessment.

3.2.7 Anterior-posterior variation

One hypothesis that we had based on prior literature was that there may be anterior-posterior differences in some aspects of hippocampal structure (Poppenk *et al.*, 2013; Strange *et al.*, 2014; Plachti *et al.*, 2019). We thus plotted select features of interest across the anterior-posterior axis within each subfield. All features can be seen in Appendix 4, where we additionally fit linear trends to the data to determine whether anterior-posterior gradients were present in any subfield. In Figure 18 we display the features mean neuronal density (Mean(y)), thickness, and gyrification that most clearly differed between subfields and are of immediate interest in MRI.

3.3 Results

3.3.1 Manual tracing

Figure 13 shows BigBrain coronal slices alongside manually segmented subfields in the head, body, and tail of the hippocampus, as well as corresponding 3D models. Several features were detected in tracings of the hippocampus in BigBrain that were not detected in previous *in-vivo* MRI work that we know of. Clusters of pyramidal cells or 'islands' can be seen on the inner surface of the subiculum (stratum lacunosum), which have been observed in histology throughout the presubiculum (Duvernoy, Cattin and Risold, 2013; Ding and Van Hoesen, 2015). A medial and anterior fold along the vertical component of the uncus, approximately 0.3 mm thick and up to 3.6 mm in length, was observed, as described in (Duvernoy, Cattin and Risold, 2013; Ding and Van Hoesen, 2015). Finally, numerous gyrifications throughout the posterior body and tail of the hippocampus were observed, which have previously also been observed using MRI in (Chang *et al.*, 2018), although not to the extent seen here. This was most prominent in CA1 but was also present in the DG and in CA4, which followed the same gyrification scheme as CA1.

Models of the dentate gyrus alone and additional anatomical notes can be found in Appendix 4. Total volumes of each subfield can be seen in Table 1. Note that these volumes are smaller than what is typically reported in MRI. This may be due to our exclusion of alveus and SRLM laminae, which can be hard to differentiate from partial voluming in MRI but may also be influenced by tissue shrinkage during histological processing. Furthermore, the issue of partial voluming in MRI may be exacerbated by the presence of gyrifications, which appeared more prominent in the right BigBrain hippocampus. These gyri are discussed in greater detail in Appendix 4.

Table 1. Volumes of each manually defined subfield (mm³).

	Left	Right
Sub	345.9	282.5
CA1	574.0	534.0
CA2	46.6	40.6
CA3	66.9	54.4
CA4	109.0	107.4
DG	140.1	131.1

3.3.2 Topological unfolding

Figure 14A shows the proximal-distal and anterior-posterior Laplacian solutions that make up the two axes of our topological unfolded space. The dentate gyrus (DG) was not unfolded. Although it was easily distinguishable from other subfields by its very high cell density it is topologically disconnected from the rest of the archicortex, and therefore would be out-of-plane (i.e. perpendicular) to our unfolded space (see Figure 13 for visualization). Figure 14B shows a mid-surface mesh of the hippocampus, coloured according to manually segmented subfields as in Figure 13. This surface was then mapped to 2D unfolded space according to the anterior-posterior and proximal-distal Laplace solutions. In unfolded space, subfields are relatively constant from anterior to posterior, with subiculum being proportionally larger in the very anterior and smaller in the very posterior extent. However, these differences may be artifacts of manual segmentation since these regions are very small in native space. This unfolding is illustrated in our online video (created through linear interpolation of all points between native and unfolded space).

3.3.3 Characterization of the hippocampus in unfolded space

Figure 15 shows a full characterization of the left and right BigBrain hippocampi with respect to the 5 morphological and 10 laminar features. These features are illustrated at the top of the figure, but additional details can be found in Chapters 3.2.4 and 3.2.5. As in related work (Duvernoy, Cattin and Risold, 2013; DeKraker *et al.*, 2018), thickness was highest in the subiculum and CA4 and lowest in CA2. Curvature was generally high in subiculum, which reflects its outward curling away from the rest of the hippocampus. In CA1, vertical bands of positive and negative values can be seen that correspond to the hippocampal gyrifications displayed in Figure 13. This region is also highlighted by our gyrification measure, which differs from curvature in that it does not vary by direction. Inner surface texture shows an almost honeycomb texture that is most prominent in the subiculum, where subicular ‘islands’ of neurons are found in stratum lacunosum (Duvernoy, Cattin and Risold, 2013). Outer surface texture appears smoother, and more closely resembles the mid-surface curvature measure. Note that the surface textures measures differ from the curvature measure only in that they capture very local details. Thus, they may not be available in lower resolution data. By contrast, features such as thickness and gyrification may be especially of interest in translation of this work to MRI, particularly because they show such clear distinction between subfields.

Of the laminar features computed here, Mean(y) was highest in region CA2, which also agrees with the high neuronal densities observed in this region (Duvernoy, Cattin and Risold, 2013). Mean(x) showed almost the inverse pattern, with high values in all regions except CA2. This means that the distribution of neurons was shifted towards the inner surface in CA2. SD(x) was highest in CA2, indicating a wide distribution of neurons relative to the thickness of that tissue. This was counter-intuitive since in native space CA2 appears to have a tight distribution of neurons; however, relative to its small thickness the distribution is wide. The remaining 8 laminar features become more complex and quite similar to Mean(y), Mean(x) or SD(x). Thus, some of these features may be redundant. We nevertheless included them for consistency with previous work in the neocortex (Amunts *et al.*, 1999). Although we did not perform any systematic

comparison, there is visibly very high consistency in all features between the left and right hippocampus, particularly with respect to the subfields.

3.3.4 Unsupervised identification of hippocampal subfields using combination of morphological and laminar features

By visual inspection, many of the features in Figure 15 show a clear distinction between the different manually defined subfields. Therefore, we sought to determine whether a combination of these features could be used to derive some or all of the subfield boundaries between subiculum and CA1 to CA4 computationally, using PCA followed by k-means clustering (see Materials and Methods for details). In this endeavour we also examined whether morphological or laminar features in isolation would be sufficient to allow for successful clustering, i.e. to derive clusters that closely resemble gold standard hippocampal subfields. For consideration of morphological features, we excluded surface textures given that they include subicular ‘islands’, which arguably also qualify as laminar features (see Appendix 7, and (Ding, 2013; Ding and Van Hoesen, 2015) for further discussion). Figure 16 shows the results of unsupervised clustering of the combined feature sets, laminar features only, and morphological features only. We compared clusters to their closest corresponding manually defined subfield (gold standard) using DSI in Table 2. When all features were combined in this analysis, good (0.7) to very good (0.8+) overlap was found for most subfields. Specifically, subfields subiculum, CA1, as well as combined CA2 and CA3 showed overlap with gold standard segmentations. Manually defined region CA2 had two clusters that overlapped with it (orange and green in Figure 16). The green cluster corresponded to the densest regions of CA2 (e.g. where Mean(y) and SD(x) were high), and several other laminar features echoed this pattern. The fact that multiple features showed this pattern may have contributed to the generation of two clusters in CA2. In other words, the variance within CA2 may have been amplified by the presence of redundant features. Using a combination of labels CA2 and CA3, as is often done in MRI segmentation protocols (Yushkevich *et al.*, 2015), increased the DSI as expected. We note that subfield CA4 did not emerge as a unique cluster and was instead included in the same cluster as CA1 or CA3. This remained true even when the number of clusters (k) was increased up to k = 16

(Appendix 7). Overlap of CA4 with CA3 is to be expected given their topological closeness but overlap with CA1 is more surprising. One possible explanation is that despite their topological separation, both of these regions were thicker, had higher gyrification, and contained a lower density of neurons than CA2 and CA3 (see Figure 15; CA4 is at the very top of each map). Relabeling clusters 1 and 2 (when they were present) past a proximal-distal distance of 200 allowed us to force a separation based on its break in continuity (i.e., separation in unfolded space, see Table 2). It should be noted that this latter approach is not purely data-driven and only offers a heuristic that is built on *a-priori* knowledge. Finally, the current analyses did not reveal any evidence for the subregions of the subicular complex as described by (Ding, 2013). This is not surprising because BigBrain only contains a single contrast (neuronal cell bodies); other contrasts (particularly myelin) or even immunochemical profiles are typically used to detect these subregions (Ding, 2013; Ding and Van Hoesen, 2015). Converging evidence was obtained for these results using different numbers of clusters, k , in k -means clustering, and using a different clustering algorithm, i.e., hierarchical clustering (see Appendix 7).

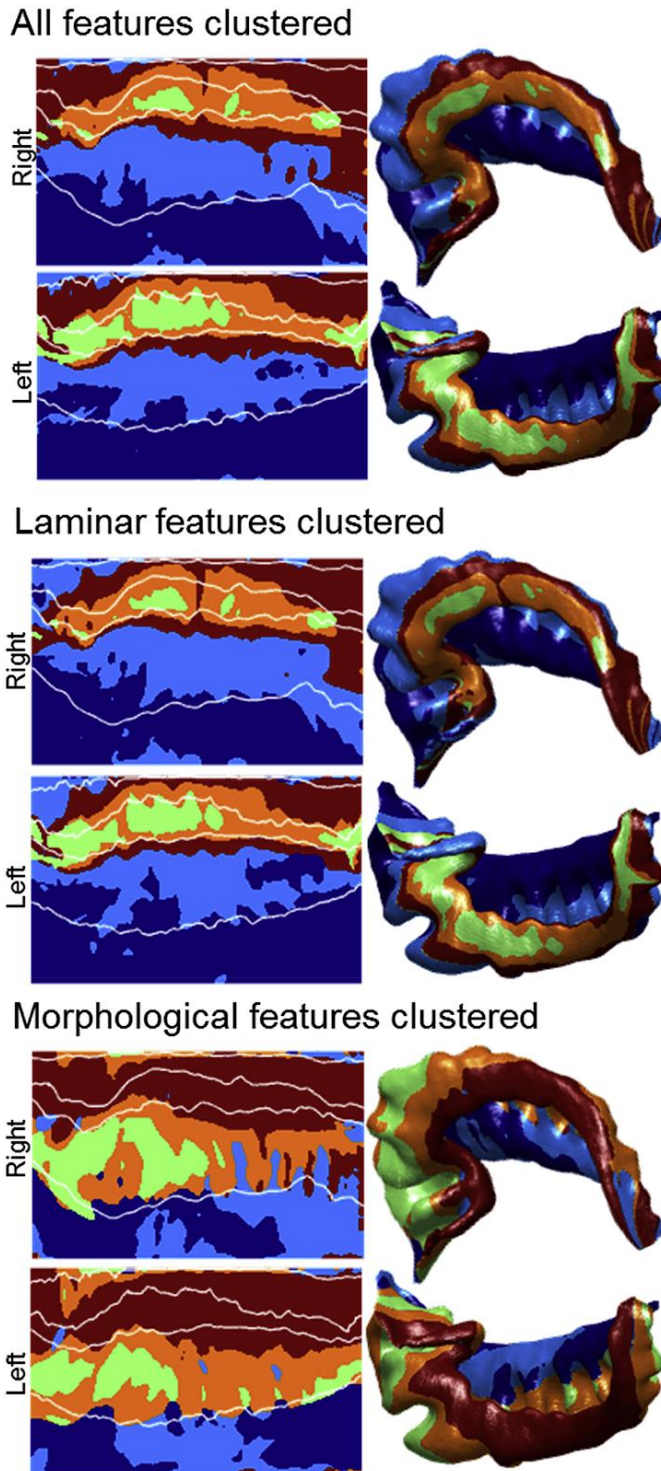


Figure 16. Unsupervised k-means clustering of features. The left images show k-means clusters in unfolded space at $k = 5$, with manually defined subfield borders overlaid in white. The right images show the same data in native space, with 10% anterior and posterior edges extrapolated by nearest neighbour. Clustering was

completed for the combined set of all features, laminar features only, and morphological features only.




Figure 17. Exploration of inherent feature variance. A) shows PCA component loadings from each feature with a dotted line at 1% after which subsequent components were discarded. B) shows the feature loadings of the first 8 components, with multiple rows for the various smoothing kernels applied to each feature. C) shows the correlation between all features, with separate boxes around morphological and laminar features. D) shows a visualization of the first 8 components, with manually defined subfield borders overlaid in white.




Figure 18. Features of interest plotted with respect to the anterior-posterior axis of the hippocampus. Colours indicate manually defined subfields, and shaded areas indicate standard deviation. Data are combined across the left and right hippocampi. a.u. stands for arbitrary units, see Chapter 3.2.7 for additional details.

Table 2. DSI between k-means clusters and their closest corresponding manually defined subfield.

	All features		Laminar features		Morphological features	
	Left	Right	Left	Right	Left	Right
Sub	0.87	0.83	0.83	0.79	0.92	0.87
CA1	0.67	0.63	0.61	0.61	0.80	0.74
CA2	0.72	0.67	0.73	0.68	0	0
CA3	0.56	0.55	0.53	0.52	0.59	0.56
CA2&CA3	0.84	0.77	0.84	0.75	0.84	0.80
CA4	0.59 ^a	0.41 ^a	0	0	0	0

^aAfter additional post-processing (see text).

3.3.5 Unsupervised identification of hippocampal subfields using morphological or laminar features in isolation

We next asked whether subsets of features (i.e. morphological features alone or laminar features alone) could be used to derive hippocampal subfield borders. Laminar features alone were able to capture most boundaries with good accuracy, with the exceptions of CA1, CA2, and CA3 which had DSI below 0.7. (Figure 16; Table 2). Again, combining CA2 and CA3 lead to good (0.7+) agreement with manually defined gold standard segmentations. CA1 was less well defined using only laminar features, and indeed there is some disagreement over the exact border between subiculum and CA1 in the histological literature (some disagreement may depend on the inclusion of prosubiculum as its own region or simply a transition zone; see (Wisse *et al.*, 2017)). Morphological features alone revealed two clusters within subiculum and two within CA1 and did not differentiate between CA2 and CA3 at all. Clustering using these features also highlighted boundaries surrounding CA4, but CA4 did not contain a unique cluster. Rather, the same clusters that were assigned to CA1 were assigned to CA4, similarly to when all features were included in clustering. However, it is worth noting that when their topological separation is considered visually, CA4 can easily be distinguished from CA1. Overall, with the exception of differentiating CA2 from CA3, morphological features were sufficient to delineate hippocampal subfields with very good (0.8 + in most cases) accuracies, at a level similar to clustering based on the combination of all features.

3.3.6 Relative contributions of individual features to subfield clustering

In order to better understand the inherent structure of the data used in the above k-means clustering of all features, we revisited the PCA that guided clustering and examined various PCA metrics. Figure 17A shows the total variance explained by each PCA component; only the first 8 components explained >1% of the variance and were included in subsequent analyses. Figure 17B and D shows a breakdown of the first 8 principal components. The first and most prominent component was most highly correlated with most laminar features, except Mean(x), Skew(x), and Mean(x.d) which showed an anti-correlation. Visualization of this component shows consistently high values in CA2. This makes sense since most laminar features showed uniquely high values in CA2, while Mean(x), Skew(x), and Mean(x.d) contained low values in this region (Figure 15). Subsequent components explain a decreasing portion of the total variance in the data, but correlate with different input features. Visual inspection of these components shows that some loosely follow the contours of the subfields. For example, component 3 quite clearly alternates high and low between subiculum, CA1, CA2, and CA3. Others, particularly components 5–8, appear to contain little subfield-related variance and may reflect noise captured by the later components. Interestingly, components 2 and 3 appear to show gradual anterior-posterior differences, with higher values in the anterior and lower values in the posterior extent in component 2, and the opposite pattern in component 3.

Of the features used in this analysis, some were more correlated with each other than others (Figure 17C). In particular, all morphological features tended to be correlated with each other. All laminar features tended to be correlated or anti-correlated with each other, with only small correlations between morphological and laminar features. The fact that laminar features tended to be uncorrelated with morphological features is in line with the goal of the Equivolume model (Waehnert *et al.*, 2014), which we applied in order to remove the effects of curvature on laminar displacement. Thus, overall, when modelled in 3D using the appropriate methods, morphological and laminar features represent different levels of structural information about tissue within the hippocampus. One

notable exception is that thickness was anti-correlated with many laminar features. This could be due to measurement bias: If thickness were overestimated in manual segmentation, then laminar profiles would be systematically shifted. However, this may also relate to how microstructural differentiation of local tissue features contribute to overall cortical thickness (see Wagstyl *et al.*, 2018 for discussion).

3.3.7 Structural variation along the longitudinal hippocampal axis

In a final set of analyses, we aimed to explore qualitatively whether subfields would show differences in feature composition along the anterior-posterior axis of the hippocampus. Towards this end we visualized possible trends along the axis in each manually-defined subfield (Figure 18). We primarily focused on the features gyrification, thickness, and mean neuronal density (Mean(y)), given that these features showed high contrast between different subfields. (Data for all other features are included in Appendix 4). Note that with this visualization, a high degree of separation can be seen between some subfields, as previously described (see Figure 15). Thickness and gyrification tended to show lower values at the anterior and posterior extremes, or in the vertical component of the uncus and tail of the hippocampus. This pattern was also observed during manual tracing (Figure 13). However, in the remainder of the hippocampus, namely the head and body, thickness remained relatively constant in each subfield while gyrification gradually decreased, as observed during manual tracing. This is most notable in CA1 where gyrification is especially prominent (note that in the Appendix 4, we also report a similar linear decrease in gyrification in CA3, but at a much smaller scale). Neuronal density was notably lower in most subfields in the anterior sections, approximately corresponding to the vertical component of the uncus. Additionally, subfield CA1 and CA4 showed linear increases in density from anterior to posterior (Appendix 4). Overall, these visualizations suggest that anterior-posterior differences are clearly present in gyrification in CA1, and in density in CA1 and CA4.

3.4 Discussion

In the present study we show, for the first time, unsupervised clustering of human hippocampal subfields that closely resembles the manually defined gold standard. We

additionally show that morphological features alone are sufficient to derive most hippocampal subfield boundaries. Moreover, our findings reveal that some features, most notably gyrification in CA1, showed within-subfield differences along the anterior-posterior hippocampal axis. The current study sheds new light on the relationship between hippocampal topology, morphology, and laminar cytoarchitecture with respect to hippocampal subfields and the anterior-posterior axis.

3.4.1 Structural characterization of the hippocampus in BigBrain

Manual tracing and 3D modelling of the hippocampus (Figure 13) at the level of resolution available in BigBrain revealed several features not seen in any 3D atlas that we are aware of. First, medial folding in the posterior end of the vertical component of the uncus was observed, similar to the inward ‘curling’ of the CA fields around the innermost DG in the rest of the hippocampus. Second, ‘islands’ of pyramidal neurons were present in stratum lacunosum in the subiculum. Third, gyrifications were present throughout the head, body, and tail of the hippocampus but were most prominent in CA1. These gyrifications were also echoed in the underlying DG (where the term dentation is often used to refer to this feature), and region CA4 that the DG partially encircles. Each of these features has been described in histology (Duvernoy *et al.*, 2013; Ding and Van Hoesen, 2015), but has not been reconstructed in a 3D model at this level of detail. For example (Adler *et al.*, 2018), and (Iglesias *et al.*, 2015) both performed detailed and fully 3D segmentation of the hippocampus and its subfields using *ex-vivo* MRI data, with additional histological data in the same participants provided by Adler *et al.* Our approach extends these studies by utilizing higher-resolution tracing and by using histological cues inherent in the same images. Furthermore, our manual traces and quantitative analyses fully respect the topology of the hippocampus and, in turn, the continuity of each subfield throughout the entire length of the hippocampus. We note that the topology developed here does not cover the dentate gyrus, which has its own topological arrangement that is perpendicular to the rest of the cortex (including archi- and neo-cortex). This difference in topology arises from a different trajectory in ontogeny, in which the DG ‘breaks with’ the rest of the cortex and wraps around the distal-most archicortex, i.e., CA4 (Duvernoy *et al.*, 2013; Nieuwenhuys *et al.*, 2013). In

future work, the DG could also be unfolded using a general framework similar to what is presented here. Critically, however, this approach would require employment of endpoints in a different plane.

After applying our topological unfolding framework, we computationally extracted morphological and laminar features from the hippocampus (Figure 15). Many of these features agree with qualitative descriptions by neuroanatomists, as discussed in Chapter 3.3.3. Some of these features may be informative for *in-vivo* imaging as well. For example, measures of thickness and gyrification can be obtained under our topological unfolding framework given sufficiently detailed segmentations, regardless of the availability of cytoarchitectonic features. These two features in particular show good contrast between subfields subiculum, CA1, and CA4. Thus, they could be explicitly leveraged to guide segmentation or registration to histological reference materials in future MRI work. This may have been underappreciated in other *in-vivo* studies, including our own previous MRI study, where some of the gyrifications in the body and tail of the hippocampus could not be detected. This prior lack of detail would also lead to overinflated thickness measures, larger overall volumes, and perhaps differences in the proportional sizes of some subfields along the anterior-posterior extent of the hippocampus. Quantitative MR, such as T1 mapping, may additionally provide cues to approximate cytoarchitectonic features. Indeed, in our previous work we observed higher T1-weights in CA2 and CA3 (DeKraker *et al.*, 2018), which may be driven by the higher neuronal cell densities observed in the current study. Thus, the features described here show clear promise for characterizing or segmenting the hippocampus in future MRI work.

3.4.2 Unsupervised clustering of all features reveals hippocampal subfields

We performed unsupervised clustering of all features to determine whether we could identify the classically described hippocampal subfields using a completely unsupervised computational approach. Results from clustering yielded generally high overlap with manual subfield segmentations in most cases (Table 2), with several exceptions that are outlined in Chapter 3.3.4. One particularly interesting observation was that CA4 was

consistently assigned the same cluster as CA3 or CA1, even though it shares no topological boundary with CA1. The shared structural elements between CA4 and CA1, particularly their relatively high thickness, gyrification, and low density of neurons, may relate to why certain diseases, such as subtypes of epilepsy, selectively affect CA1 and CA4 similarly (Duvernoy *et al.*, 2013; Blümcke *et al.*, 2013). In future imaging work, CA1 and CA4 may be differentiated from each other, particularly under our unfolding framework, due to their topological separation.

Finally, to further explore the inherent dimensionality of the structural feature space under consideration, we also examined its principle components (Figure 17). In these analyses, we noted that the most prominent components varied in ways that followed the contours of some or all subfield borders (see Chapter 3.3.4). This pattern suggests that the inherent structural variance in the hippocampus most naturally follows a proximal-distal patterning, as seen in the classic histological subfield definitions. Some components additionally hinted at inherent anterior-posterior differences across the hippocampus.

Note that the analyses performed here assume discrete subfield boundaries, but some recent literature considers select regions simply as transition areas (e.g. prosubiculum could be considered a discrete region or a transition area between CA1 and subiculum, see Olsen *et al.*, 2019). In Appendix 7, alternative clustering results consistently show a boundary in this region, but, even with large numbers of clusters in k-means clustering, a discrete prosubiculum region doesn't clearly emerge. This most closely aligns with the idea that prosubiculum is a transitional area, showing graded differences between CA1 and subiculum rather than its own unique characteristics. However, it should be noted that we did not examine myelo- or chemo-architecture, which may indeed differentiate prosubiculum from its neighbours (Ding and Van Hoesen, 2015).

3.4.3 Morphological features are sufficient to approximate most subfield boundaries

In addition to clustering using all features, we also asked whether hippocampal subfields could be derived using only the subset of morphological features or the subset of laminar features (Figure 16 and Table 2). Clustering using laminar features revealed all

hippocampal subfields with reasonable accuracy with respect to manually defined gold standard segmentations (except CA4, and CA2 versus CA3 were not easily distinguished from one another). This was expected, since laminar features provide the key criteria used by histologists to define subfield boundaries (e.g. Duvernoy *et al.*, 2013; Ding and Van Hoesen, 2015). However, when we examined morphological features alone, we also found unsupervised clusters that closely resembled subfields subiculum, CA1, and a combined CA2 and CA3 region. Additionally, CA4 was assigned the same clusters as CA1, similar to when clustering was performed on all features combined. This outcome was not expected based on histological data and provides support for the notion that morphological features capture an independent set of subfield-related structural elements. The observation that morphological features are sufficient to determine most subfield boundaries holds great promise for future refinement of MRI protocols for subfield delineation, given that histological- or laminar-level details are not available in current imaging protocols. Indeed, many of the MR-based subfield segmentation protocols presently available rely on some combination of structural landmarks within or surrounding the hippocampus, but only indirectly on morphological features (see Yushkevich *et al.*, 2015a). However, some caution should still be exercised here as the current results were obtained from a single individual.

3.4.4 Anterior-posterior structural variation

Anterior-posterior structural differences in the hippocampus are particularly of interest, given the growing body of literature suggesting functional gradients along the longitudinal axis of the hippocampus (e.g. Strange *et al.*, 2014; Poppenk *et al.*, 2013; Plachti *et al.*, 2019; Zeidman and Maguire, 2016). Structural anterior-posterior gradients are difficult to assess using conventional histology, given that coronal or sagittal sections are typically out-of-plane with respect to the different subfields in most of the hippocampal head and tail. This highlights the utility of the 3D BigBrain dataset. Figure 18 shows features gyrification, thickness, and neuronal density along the anterior-posterior axis of the hippocampus. Most notable anterior-posterior differences included differences in most features at the very anterior and posterior extents of the hippocampus. Previous work by (Ding and Van Hoesen, 2015) described the anterior most region - the

vertical component of the uncus - as containing modified subfields that were much thinner than their counterparts throughout the rest of the hippocampus, consistent with our observations.

Gyrification was low in the anterior uncus, high in the remainder of the hippocampal head, and gradually decreased towards the posterior end of the hippocampus, most notably in CA1. Qualitatively, similar trends in gyrification have been observed in our previous 7-Tesla MRI study (DeKraker *et al.*, 2018) and in other work (Chang *et al.*, 2018). However, both of these studies were limited in their ability to detect small gyrifications (i.e. those detected in this study had peak-to-peak distances as low as 2mm). Biophysical models of the development of gyrification suggest a relationship between gyrus size and cortical thickness (Zilles *et al.*, 2013; Striedter *et al.*, 2015). Yet, no systematic anterior-posterior differences in thickness were seen in the present data despite clear decreases in gyrification size towards the posterior extent. Other structures such as white matter might also constrain gyrification patterns (Striedter *et al.*, 2015), which may additionally have consequences for functional properties of different gyri. For example (Henderson and Robinson, 2014), examined gyrification and structural connectivity in the neocortex and found more unified or modular graph theoretical properties within gyri, as opposed to sulcal regions which were more diffusely connected or hub-like. Similarly (Plachti *et al.*, 2019; Libby *et al.*, 2012), recently performed parcellation of the hippocampus according to its functional connectivity and observed divisions primarily along the anterior-posterior extent of the hippocampus, rather than across subfields (although some proximal-distal clustering was also observed, as in the present study). This functional parcellation may even relate to modular divisions of function within a given gyrus, as proposed for the neocortex by (Henderson and Robinson, 2014).

It is also interesting that neuronal density increased from anterior to posterior sections in subfields CA1 and CA4 in the present study. It should be noted that the current methods cannot differentiate density from neuronal size, but other related work has also found similar effects in density (Dam and Mouritzen Dam, 1979).

3.4.5 MRI applications and future directions

There are several clear implications of this work for *in-vivo* neuroimaging studies of the human hippocampus. Firstly, as described in the Chapters 1.4 and 3.1, considerable research has gone into defining the hippocampal subfields according to available landmarks in MRI. Yet, relatively few studies have explicitly investigated the 3D shape of the underlying archicortex in which those subfields are embedded. Most subfield delineations rely directly or indirectly on manual segmentation performed predominantly in the coronal view, in which the only visible gyrifications are in the hippocampal head (Yushkevich *et al.*, 2015a). Because of this complication, many protocols have elected not to divide subfields in this region, or else to simplify them. The same concern also applies in the hippocampal tail, albeit at an even finer spatial scale. In our previous 7-Tesla MRI study (DeKraker *et al.*, 2018), we demonstrated the use of topological modelling to overcome this challenge, but we note that this previous work captured fewer details than the current histological study due to resolution limitations. In future work, we anticipate it will be possible to explicitly model gyrifications and all other aspects of hippocampal topology considered in the present study. In addition, topological modelling of the hippocampus in MRI can open new lines of structural investigation, including the development of more precise measures of thickness, gyrification, and, at higher resolutions, potentially laminar features. Our approach promises to be of particular value in overcoming systematic partial voluming with white matter structures within the hippocampus and in adjacent structures. At a broader level, consideration of the hippocampal features highlighted in our study may also allow researchers and clinicians to link pathological changes, or behavioural and cognitive phenotypes, to more specific structural elements within the hippocampus. Finally, our current results may also help bridge the gap between microcircuit level modelling of the hippocampus and subfield modelling at the scales currently available in MRI. Because we found that subfields are characterized by unique morphological feature combinations, we may even be able to infer a relationship between subfield microcircuitry and fMRI measures in future work.

3.4.6 Data and resources made available

Alongside this publication, we released our detailed manually defined hippocampal subfields, unsupervised clustering results, topological unfolding framework, Equivolume laminar model solutions, and each of the unfolded morphological and laminar features computed here for the BigBrain dataset. These resources can be used as templates in other studies. Alternatively, registration of these features to new data in our unfolded space can be used to guide future subfield segmentation. In addition, we have also made the code used in this project available via Open Science Framework (<https://osf.io/x542s/>). A toolbox for performing hippocampal unfolding, feature extraction, and other useful operations on more general datasets can be found at <https://github.com/jordandekraker/Hippunfolding>.

3.4.7 Conclusions

In the current project, we mapped the human hippocampus in detail by combining three methods. First, we used a unique dataset, BigBrain, that contains both histological-level detail and macroscopic 3D spatial context. Second, we imposed a topological unfolding framework to the hippocampus. Third, with this framework we extracted a set of morphological and laminar features, the latter of which have been used prolifically in neocortical characterization and parcellation. Using these methods, we highlight three novel empirical observations. First, unsupervised clustering of these features closely resembles classically defined hippocampal subfields. Secondly, despite traditional reliance on laminar features in histology, morphological features alone are sufficient to closely approximate most hippocampal subfields. Finally, some features such as gyrification in CA1 show, at least qualitatively, subfield-specific anterior-posterior differences that might relate to functional differences described in the extant literature. Overall, these findings highlight new structural characteristics of the hippocampus, and offer promising avenues for improved delineation and characterization of hippocampal subfields using *in-vivo* neuroimaging.

Chapter 4

4 Hippocampal Automated Topology

This Chapter focuses on the automation and generalization of the hippocampal coordinate framework and feature extraction methods discussed Chapters 2 and 3 to new datasets. This includes the translation of subfield boundaries defined using 3D histology in Chapter 3 to new MRI data via our unfolded (or surface-based) coordinate framework. This framework requires classification of hippocampal tissues into grey matter, high myelin strata (or SRLM), and surrounding structures. In particular, detailed separation of different hippocampal folds or digitations via the SRLM on the inner surface, and surrounding alveus, white matter, or CSF on the outer surface is required. This is challenging to do using traditional segmentation methods since these structures are often thin (e.g. as low as 0.3mm in thickness) and their location can be variable depending on the folding configuration of each subject's hippocampus. To overcome these challenges, we pooled data from previous projects and used it to train a deep neural network (U-Net) for tissue classification. We then applied incremental learning to further train U-Net on the Human Connectome Project (HCP1200) dataset, specifically on 0.7mm³ isotropic T2w images. With post-processing of resulting tissue classifications via a template shape injection method, we were able to apply the unfolded coordinate framework to all subjects in the HCP1200 dataset. We then applied the subfield boundaries derived in Chapter 3 to all subjects in unfolded space. Collectively, we refer to these methods as Hippocampal Automated Topology (HAT). Features and subfields extracted or applied with these methods were compared to previous work and other extant automated hippocampal subfields segmentation methods. Overall, HAT yielded structural features that agree with previous work and *ex-vivo* studies, and subfields defined using this method compared favourably to other extant methods in terms of detail and qualitative comparison to anatomical literature.

4.1 Introduction

State-of-the art methods for neocortical parcellation make use of surface-based representations. This facilitates feature extraction (e.g. thickness, gyrification index, or

quantitative MRI), but also allows similar parcellation schemes to be applied to many subjects despite broad differences in gyral and sulcal patterning. This is challenging to do using more traditional 3D registration based methods where topological breaks often occur when registering subjects with different fine scale gyral or sulcal patterns. A surface-based approach can help overcome this problem by projecting cortical tissue to an unfolded flat or spherical space with 2D topology (see Dale, Fischl and Sereno, 1999; Fischl, Sereno and Dale, 1999; MacDonald *et al.*, 2000; Zijdenbos, Forghani and Evans, 2002; Kim *et al.*, 2005; Glasser *et al.*, 2016; Research and Case Medical Research, 2019). Parcellation performed in this space is constrained from 3D to 2D and can subsequently be projected back to each subject's native space despite differences in folding patterns. In recent work, we developed a surface-based (or topological) framework for examination of the archicortex, specifically, the hippocampus. As in the neocortex, this enables detailed feature extraction and helps overcome broad inter-individual variability in gyrification, or digitation as it is often referred to in the hippocampus. This is especially critical given recent reports of inter-individual morphological variability in the hippocampus and its subfields (Ding and Van Hoesen, 2015; Cai *et al.*, 2019; de Flores *et al.*, 2019; DeKraker *et al.*, 2020). One major challenge with this surface-based approach, which we aim to address here, is the need for careful manual delineation of tissues that separate the folds or digitations of the hippocampus. This process is costly in time and expertise and reduces the reproducibility and scalability of this method.

Briefly, our surface-based approach to hippocampal feature extraction and parcellation, or subfield delineation, involves indexing hippocampal grey matter on three geodesic axes: anterior-posterior, proximal-distal, and laminar (or inner-outer). The anterior-posterior and proximal-distal indices make up a 2D topologically organized space that can readily be flatmapped. This indexing is determined according to the Laplace equation, where hippocampal grey matter corresponds to the domain and anatomically motivated structures at each edge (e.g. the hippocampal-amygdalar transition area at the anterior terminus) make up boundary conditions. The hippocampal sulcus and surrounding high myelin laminae (strata radiatum, lacunosum, and moleculare or SRLM) separate the inward curling of the hippocampus, as well as the inner boundaries between digitations, while the alveus and surrounding structures (including the third ventricle and

inferior longitudinal fasciculus) separate the outer boundaries between digitations. Topological breaks or bridges between digitations in this grey matter label can introduce major distortions to the Laplacian solution, and so highly detailed segmentation of these structures is required.

Artificial neural networks, and particularly the U-Net architecture (Chen *et al.*, 2018) are becoming increasingly popular in medical image segmentation. U-Net can take advantage of local and global image features in 3D, and in principle it could also discover 2D topological structure in 3D data or leverage subtle differences in thickness or intensity between hippocampal subfields. However, there is no evidence that current applications of U-Net for hippocampal and hippocampal subfield segmentation leverage such sophisticated topological modelling or subtle subfield-related intensity differences. Instead, results closely resemble the manually labelled data which they were trained on, which does not show detailed digitations and corresponding topological shifts in subfields (Shi, Cheng and Liu, 2019; Zhu *et al.*, 2019; Yang *et al.*, 2020). This highlights a need for more elaborate methods and/or more detailed training data.

In the current work, we apply U-Net segmentation in concert with our previously developed hippocampal unfolding framework, which we jointly refer to here as Hippocampal Automated Topology or HAT. Starting with detailed hippocampal segmentations from a collective of previous studies, we aimed to develop an automated method for obtaining detailed segmentation of hippocampal tissue classes in the Human Connectome Project (HCP1200) dataset (Glasser *et al.*, 2013). These tissue classes can then be used to unfold hippocampal grey matter under our previously developed hippocampal coordinate framework (DeKraker *et al.*, 2018). That is, we use U-Net to segment grey and white matter tissue classes within the hippocampus (most critically hippocampal grey matter and the high myelin SRLM, with background tissues spanning CSF, surrounding white matter, and alveus). Following post-processing, we then apply our previously developed unfolded coordinate framework and define hippocampal subfields according to highly detailed boundaries derived from 3D histology in previous work (DeKraker *et al.*, 2020). It should be noted that while the same boundaries are applied to all hippocampi in unfolded space, they may vary considerably between

subjects in native space depending on the folding configuration (particularly the number and prominence of digitations) in each subject's hippocampus. This entire pipeline is overviewed in Figure 19.




Figure 19. Overview of full Hippocampal Automated Topology pipeline. Step 1: left and right hippocampi are cropped and resampled to 0.3mm^3 isotropic obliquely to the hippocampal long-axis. Step 2: images are segmented via U-Net architecture and then post-processed using template shape injection. Step 3: Previously developed Laplace coordinate framework is applied to the domain of hippocampal grey matter tissue. Step 4: A single subfield atlas defined in 2D unfolded space, in this case generated from BigBrain 3D histology, is propagated to a given subject's native hippocampal folding configuration. Shown here is the example of the right hippocampus of HCP1200 subject 108020.

4.2 Methods and data

4.2.1 Preprocessing and resampling

Initial manually labelled data was reused from a variety of different studies, including 20 samples (from 10 subjects) from the HCP1200 dataset (Glasser *et al.*, 2013; Ferko *et al.*, 2017), 22 samples (from 16 subjects) from an open source UPenn *ex-vivo* MRI dataset (UPennExVivo) (Adler *et al.* 2018) with manual adjustments to suit our unfolding protocol, and 24 samples (from 12 subjects) used in our initial hippocampal unfolding protocol development (UWO7T) (Jordan DeKraker *et al.* 2018). Using incremental learning, we were able to achieve highly detailed semi-manual segmentations of left and right hippocampi from the full HCP1200 dataset.

Prior to U-Net segmentation, the left and right hippocampi were extracted from each whole-brain T2w scan, then each was rotated to be coronally oblique to the long-axis of the hippocampus, resampled to 0.3mm^3 , and cropped to a $128 \times 256 \times 128$ voxel volume surrounding the hippocampus. This was performed by computing an affine transformation using FSL's FLIRT (Jenkinson and Smith, 2001; Jenkinson *et al.*, 2002) to a single group-averaged atlas (CITI168 (Tyszka, Michael Tyszka and Pauli, 2016) which was combined with a manually defined rotation to coronal oblique and a manually defined left and right cropping. These steps (linear registration to the CITI168 atlas, rotation to coronal oblique, and cropping) were all performed using a single linear interpolation to minimize loss of image contrast. Taking advantage of the high left-right symmetry of the hippocampus, left hippocampi were flipped sagittally to resemble right hippocampi in order to simplify segmentation.

4.2.2 U-Net implementation

In the current work, we aimed to segment tissue classes within the hippocampus including grey matter, high myelin SRLM, background tissues spanning CSF, surrounding white matter, and alveus, and extra-hippocampal structures that act as boundaries for our unfolding coordinate framework (hippocampal-amygdalar transition area, indusium griseum, medial temporal lobe cortex, and the dentate gyrus within hippocampal grey matter). To this end, we applied a specific variant of a U-Net known as

highres3Dnet (Li *et al.*, 2017), which includes design choices and hyperparameters that are able to efficiently handle large 3D images such as those often seen in medical image processing. Briefly, this includes a total of 20 convolutional layers with regularly increasing numbers of kernels and dilations, as well as batch normalizations and residual connections every two layers. The output from this network contains the same number of nodes as the input image multiplied by the number of labels. We used NiftyNet (Gibson *et al.*, 2018) software for implementation of this network, built on Tensorflow (Abadi *et al.*, 2015). During the first step of incremental learning, we experimented with several data augmentation methods and U-Net hyperparameters and selected the following configuration based on qualitative observation of resulting inferred segmentations:

- The ‘large’ variant of highres3Dnet was used, which introduces 3 additional convolutional layers.
- 64x64x64 voxel patches were uniformly sampled with 16 voxels of overlap in each direction (since segmentation quality is generally poorer at the periphery of each patch) (larger patches exceeded available GPU memory).
- Dice overlap was chosen as a loss function (Dice, 1945; Sørensen, 1948), which has also been shown to be effective in training deep neural networks for segmentation (Sudre *et al.*, 2017).
- Data augmentation included rotations of 10° in any direction, random deformations with 4 control points combined with a smoothing kernel of $\sigma=15$, and the introduction of random 3rd order bias fields in any direction of up to 5%. These were performed on all samples except the deformations which were performed on only 75% of samples.

4.2.3 U-Net training and evaluation

U-Net was trained over 4 incremental learning steps, with quality control of results between each step. When the quality of inference segmentations was good in one step, they were added to the pool of training data for the next step. Quality was assessed in two ways: any segmentation that could not be run to completion through the remainder of the unfolding pipeline was discarded, and of those that did run to completion were manually

inspected and then either discarded, manually adjusted, or added directly to the pool of training data for the next step. Though labour intensive, this protocol was exponentially more efficient than performing de novo manual segmentation of such a large dataset, while still ensuring excellent quality in the training data for subsequent incremental learning steps.

The first step in incremental learning was trained using manual segmentations collected from a variety of different studies, including 20 samples (from 10 subjects) from the HCP1200 dataset (Glasser *et al.*, 2013; Ferko *et al.*, 2017), 22 samples (from 16 subjects) from an open source UPenn *ex-vivo* MRI dataset (UPennExVivo) (Adler *et al.*, 2018) with manual adjustments to suit our unfolding protocol, and 24 samples (from 12 subjects) used in our initial hippocampal unfolding protocol development (UWO7T) (DeKraker *et al.*, 2018). The second step in incremental learning additionally included 80 samples (from 40 subjects) from in-house scans (UWO3T). The third step in incremental learning additionally included 1992 samples (from 1000 HCP1200 subjects). Finally, for the fourth step in incremental learning we decided to further tighten the inclusion criteria and include only samples for which we had very high confidence in the segmentation quality following manual inspection. This included only 1161 samples (from 763 subjects) from the HCP1200 dataset. The participant, acquisition, and preprocessing details from all of these datasets can be found in Appendix 8.

During the final incremental learning step, 15% (n=190 samples) of the data were randomly chosen and left out of the training pool for validation, and an additional 15% (n=190 samples) were left out for test benchmarking. These subsamples never split the same subjects left and right hippocampi between training, validation, and test sets which could bias results if high within-subject left-right similarity was seen. Every 1000 training iterations, the validation set was compared against manual ground truth in terms of Dice overlap using Tensorflow's 'Tensorboard' tool. Each incremental learning step was performed using 50 000 iterations, as performance on the training and validation sets showed little improvement beyond 40 000 iterations (see Figure 20A). U-Net was reinitialized with random weights between each step. Thus, steps were nearly independent, except for the fact that segmentations used in training from one step could

be generated through inference and manual correction from a previous step. Following training on the final incremental learning step, the test set was compared to manual segmentations both before and after post-processing.

4.2.4 Post-processing

In post-processing we employed a method called template shape injection (Qiu and Miller, 2008) to correct possible segmentation errors in U-Net segmentation, making labelmaps more amenable to the previously developed hippocampal unfolding methods. The basic principle of template shape injection is to perform highly fluid deformable registration of a template segmentation label map to a given subject's segmentation label map. This differs from typical registration-based segmentation methods in that the registration is optimizing labels rather than image feature similarity (i.e. registration is performed on labels encoded as integers rather than on floating point voxel intensities). Specifically, we used mean squared error between labels as the cost function, which is minimized when identical labels are overlapping. In our implementation, we apply multi-contrast SyN deformable registration using the Advanced Normalization Toolkit (ANTs) (Avants *et al.*, 2008), where each contrast is a binarized label from the U-Net segmentation outputs. The reference template that we applied was created by combining manual segmentations from the UPennExVivo dataset using a standard template building ANTs script 'buildtemplateparallel.sh' (Avants *et al.*, 2010). This template generation entails averaging all images and then registering each sample to the average, iteratively refining and sharpening the average image. The UPennExVivo dataset was selected for template building because we had high confidence in the quality of these segmentation since they contained higher resolution and contrast than other datasets while still including multiple samples.

In the current work we introduced several optional tweaks to this post-processing step, though they may be excessively convoluted and may be omitted in future versions. Currently, extra-hippocampal structures (HATA, IG, DG, MTLc, pial) were defined according to the manual protocol presented in (DeKraker *et al.*, 2018), which only defines these ROIs at their border with the hippocampus. The other boundaries of these structures were more or less arbitrary as they were not needed for our protocol, and so we

chose to exclude them from label-label registrations. Additionally, the positioning of cysts along the SRLM is somewhat spurious, and so for the purposes of label-label registration we chose to simply combine cysts with the SRLM label. Thus, only two binarized labels were used in multicontrast deformable registration of the template to a given sample: hippocampal grey matter and SRLM. We also were initially concerned in particular about loss of detail in the hippocampal SRLM, which would lead to a general 'flattening' of the digitations which it separates. To mitigate this, we performed a separate registration weighted more heavily on the binarized SRLM label to capture more of the details of this structure. We then combined this deformation with the deformation generated from the combined labels using weighted averaging (where weights were determined by distance from the SRLM such that the deformations computed uniquely from the SRLM would be preferred near the SRLM). Qualitatively we found this approach to work well, but it does introduce complications to the pipeline that may not be necessary, as noted above.

It is important to note that this post-processing strategy can introduce bias into new segmentations. Depending on the fluidity or regularization of the registration, segmentations can come to more closely resemble the reference template than the original subject-specific segmentation. In this case the reference template was generated from a group average from the UPennExVivo dataset, as described above, and so we expect some level of regression towards the mean hippocampal shape. In the case of image-image registration, high fluidity can lead to errors because of the increased size of the problem space or decreased spatial regularization. However, this was less problematic in our case given that registration was performed only between corresponding labels rather than potentially noise image voxels. To test the effect that our reference template had on final segmentation outputs, we replaced the averaged template for individual samples from the UPennExVivo dataset and qualitatively examined resulting segmentations deformed to one subject's U-Net segmentation (HCP1200 subject 105023's right hippocampus). We additionally tested different values of deformable registration parameter 'update field variance' which most closely controls fluidity.

4.2.5 Unfolding and subfield definition

The remainder of the hippocampal unfolding pipeline is exactly as described in (DeKraker *et al.*, 2018, 2020). That is, Laplace fields are calculated across the domain of the hippocampus in three axes of interest: anterior-posterior, proximal-distal, and inner-outer (or laminar). These gradients were binned into 256x128x4 points. The anterior-posterior and proximal-distal values are used as a 2D indexing system for flatmapping, or unfolding, the hippocampus. T2w was extracted along the mid-surface of the hippocampus (i.e. at half the value of the inner-outer axis) to avoid potential partial voluming with neighbouring white matter or CSF. Because of the same partial voluming concern, we did not examine laminar intensity differences as in (DeKraker *et al.*, 2020). Cortical thickness was assessed by measuring streamlines spanning the inner-outer gradient. Gyrification index was calculated as a ratio of native space surface area over unfolded surface area at each unfolded point.

Unfolded space provides intrinsic registration between subjects and so it makes up an ideal space in which to define the hippocampal subfields. In previous 3D histology work (DeKraker *et al.*, 2020) we defined these subfields in great detail and validated these definitions with unsupervised methods. Thus, we aimed to apply the same unfolded subfield definitions here. Minor differences were seen in subfield boundaries between the left and right hippocampus in this previous work, and so the two were averaged together before being applied to all HCP1200 subjects in unfolded space. Note that while these subfields are all defined the same way in unfolded space, they may take on drastically different positions when projected back to each subject's native 3D space. This illustrates the advantage of defining subfields in a topologically unfolded space rather than in native space where subfield locations can vary considerably depending on inter-individual differences in hippocampal folding or morphology. Volume of each subfield was calculated in CITI168 space in order to control for overall brain volume.

4.2.6 Comparison to other methods

Ideally, subfield definitions generated through the current method, Hippocampal Automated Topology, which we abbreviate here as HAT, could be compared to a same-

subject histological subfield definition as a ground truth, but no such data were available. Subfield definitions can also be compared to a reference material, such as 3D BigBrain which was examined in Chapter 3. This could be done in unfolded space to enforce alignment between each subject and the reference, but this comparison would be uninteresting since our HAT method already applies BigBrain subfield definitions in unfolded space. Subfields could be compared in native space using traditional alignment methods, such as 3D deformations, but 3D deformations between subjects with different hippocampal folding configurations can lead to topological breaks or gross deformations (for example, when registering a subject with 3 anterior digitations to one with 5). Thus, this method would not produce an appropriate gold standard and would instead be cruder than the topological subfield definitions implemented here using HAT. Thus, we make only qualitative observations about subfield definitions using HAT and compare them to extant anatomical literature. Additionally, we also quantitatively and qualitatively benchmark HAT against popular extant automated subfield segmentations methods.

The two most popular methods for automatically delineating hippocampal subfields in MRI are Freesurfer's (FS) hippocampal subfields extension (Iglesias *et al.*, 2015) (currently at 382 citations), and Automatic Segmentation of Hippocampal Subfields (ASHS) (Yushkevich, Pluta, *et al.*, 2015) (currently at 255 citations). We thus applied FS and ASHS to a subset of 310 HCP1200 samples (from 155 subjects) and compared them to HAT. Note that ASHS can implement different reference atlases and more recent and detailed atlases do exist. However, we chose the current version (UPenn PMC Atlas - 2016 version) for its widespread popularity. We calculated the volume of each subfield from each protocol, and then we identified equivalent labels across protocols (see Figure 22). We calculated Dice overlap between equivalent labels between each of the three protocols. Additionally, for each subject we projected subfield identities from HAT, FS, and ASHS into the unfolded space generated via HAT. We then identified the most common (mode) subfield identity at each unfolded position across all subjects in each of the protocols (Figure 24).

To further examine systematic differences between protocols, we performed label-label registration between HAT and FS and HAT and ASHS. This is illustrated for one subject

in Figure 25A, but we also wanted to look for systematic differences between protocols across all subjects, which is challenging when subjects are not aligned. Thus, we sampled each deformation field in unfolded space before averaging across subjects. For easier visualization, we generated an average mid-surface on which we then overlaid the average deformations from HAT to FS and from HAT to ASHS as vectors originating from each vertex. Applying these averaged deformations generated mid-surfaces resembling the ASHS and the FS protocols. The results from each of these operations can be seen in Figure 25B.

4.3 Results and Discussion

4.3.1 U-Net segmentation

Figure 20A shows the final step of U-Net training, though the training was monitored at all incremental learning steps, which follow a similar pattern to the final step. The Loss (1 - Dice overlap) is shown averaged across all labels for the training and validation sets over 500K training iterations. Little improvement was seen past 400K iterations. No evidence of overfitting was seen, as the validation Loss closely followed the training Loss rather than eventually increasing. It is possible that further training might result in slightly better model fits, but this would also run the risk of overfitting any possible idiosyncrasies of the training dataset.




Figure 20. Benchmarking of tissue segmentation following the final incremental U-Net training step. A) shows the training of the final U-Net incremental learning step

over 500 000 iterations. Solid colours are smoothed over iterations and faded colours indicate raw loss at each iteration. B) shows the Dice overlap between manual and automatically generated segmentations for each label of interest in the test set (15%). Results are shown both before and after post-processing via template shape injection.

We evaluated performance (Dice overlap) for each label separately in the test set after training was completed, which is shown in Figure 20B. Dice overlap scores are well above manual inter-rater scores in critical labels grey matter (GM) and SRLM, which were previously reported to be ~ 0.83 and ~ 0.71 for GM and SRLM in 7-Tesla MRI (DeKraker *et al.*, 2018), and ~ 0.79 and ~ 0.67 for GM and SRLM in a related bachelor's thesis examining application of our manually defined hippocampal unfolding in the HCP1200 dataset (Ferko *et al.*, 2017). Other labels showed lower Dice overlap but this is to be expected since they were defined mainly using heuristics rather than boundaries discernable in the images themselves. This is of little consequence for subsequent unfolding and subfield definition since extra-hippocampal labels are only used where they directly border the hippocampus.

4.3.2 Post-processing

Dice overlap scores between manual raters and automatic segmentation following post-processing via template shape injection are shown in Figure 20B. In general, slight reductions were seen for all labels following post-processing but performance remained well above manual inter-rater Dice overlap scores. As described in Chapter 4.2.3, this may reflect some regression of hippocampal morphology towards the template shape. To test this effect, we compared post-processing via template shape injection with an averaged template reference to post-processing via template shape injection with several exemplar segmentations from the UPennExVivo dataset as references. We additionally compared several different update field variance parameter values (which relates to fluidity of the registration). Results are shown in Figure 21.




Figure 21. Effect of template choice and update field variance (a parameter regulating fluidity) on template shape injection results. Only GM (red) and SRLM (orange) are shown for clarity. The original tissue segmentation contains topological breaks (green arrows) due to segmentation errors, which are problematic for subsequent unfolding. Highly fluid template shape injection from an averaged

hippocampal shape, as well as several notably different exemplars are shown. The effect of template choice is particularly notable at lower update field variance. Final parameter choices are indicated by the black box. Segmentation errors persist in this case (purple arrows), but do not include topological breaks.

We specifically chose template shape exemplars with a wide variety of overall morphologies but nevertheless, remarkable similarity was seen following injection to the original segmentation as produced by U-Net. This was especially true with higher update field variance values, which increased the overall fluidity of the deformable registrations. Update field sigmas beyond a value of approximately 15 made little difference to the final segmentation, and so we used this parameter value in the remainder of the work shown here. Note that template shape injection does not perfectly account for segmentation errors, as indicated by the purple arrows in Figure 21. The main utility of this step in the overall pipeline is that it enforces some level of topological continuity, preventing holes or bridges between adjacent structures that are not connected (e.g. bridges between digitations or across the SRLM) in order to avoid possible errors in the subsequent unfolding paradigm.

4.3.3 Unfolded subfields and features

All subjects ran to completion through our previously developed hippocampal unfolding framework, and visual inspection showed negligible distortions in the Laplace fields used to define unfolded space. Features extracted using this unfolding method are shown in Figure 22.




Figure 22. Average features across all HCP1200 subjects plotted in unfolded space. A group-averaged mid-surface in native space is shown in the top left.

Features plotted in unfolded space and averaged across all HCP1200 subjects showed remarkable similarity to previous work (DeKraker *et al.*, 2018). That is, T2w was generally high in CA1 and CA4 while being low in the other subfields. Gyrification index was high in CA1 and CA4 as in (DeKraker *et al.*, 2020), but in this case high values were also seen in distal subiculum. Thickness was high in CA1 and subiculum and low in the other subfields, especially CA2. Vertical stripes of higher thickness can be seen near the hippocampal head and tail, and also thickness values were generally higher than in previous 3D work (DeKraker *et al.*, 2020) in CA1. This may be due to the level of detail available in previous 3D histology work which allows for identification of even small digitations. When these digitations are not detectable, they may simply appear as a thicker and smoother archicortical mantle rather than a folded structure. Indeed, this previous work suggested that imprecise measurements of small digitations may conflate measures of thickness, surface area, and even T2w due to partial voluming.

4.3.4 Comparison to other methods

We aimed to compare subfield segmentations obtained under the present protocol to other popular methods Automatic Segmentation of Hippocampal Subfields (ASHS) and Freesurfer 6.0+ (FS). One example segmentation from each method can be found in

Figure 24. Average total volumes (i.e. all grey matter and SRLM labels combined) were $2933 \pm 228 \text{mm}^3$ and $2943 \pm 227 \text{mm}^3$ in the left and right FS segmentations, $2720 \pm 474 \text{mm}^3$ and $2771 \pm 482 \text{mm}^3$ in the left and right HAT segmentations, and $2622 \pm 233 \text{mm}^3$ and $2700 \pm 264 \text{mm}^3$ in the left and right ASHS segmentations. Figure 23 shows these volume differences for each subject. Note that FS includes parasubiculum, which is part of the hippocampal formation but was not included in either of the other two protocols. This may partially account for larger volumes under this protocol. We additionally note that the hippocampal head and tail were larger under this protocol, in some cases clearly exceeding the outer boundaries of the hippocampus.




Figure 23. Bland-Altman plots comparing total hippocampal volume between FS, HAT, and ASHS. Units are in mm^3 , and volumes include both hippocampal grey matter and SRLM.

We calculated the average volume of each subfield (Figure 24 second column), which varied considerably between methods. Of particular note is that the combined subiculum labels in FS are much larger than HAT or ASHS, while CA1 was considerably smaller. This may pertain particularly to the subiculum-CA1 border, which is subject to some debate in the field (see discussion by Wisse *et al.*, 2017). ASHS segmentations of CA2 and CA3 were much smaller than HAT or FS, and indeed there are many coronal slices that contain no CA2 or CA3 at all. This may in part be due to interpolation and voxel size of the reference materials: large reference voxel sizes may mean that no voxels are labelled CA2 or CA3 or, even when they are present in the reference material, they may become compressed during registration to target samples and disappear entirely. Finally, ASHS also showed CA4&DG labels that were much larger than the other two methods. This may be in part due to the inclusion of SRLM. Overall, HAT total volumes were much smaller but, proportionally, individual subfield volumes were similar or in between those obtained from FS or ASHS.

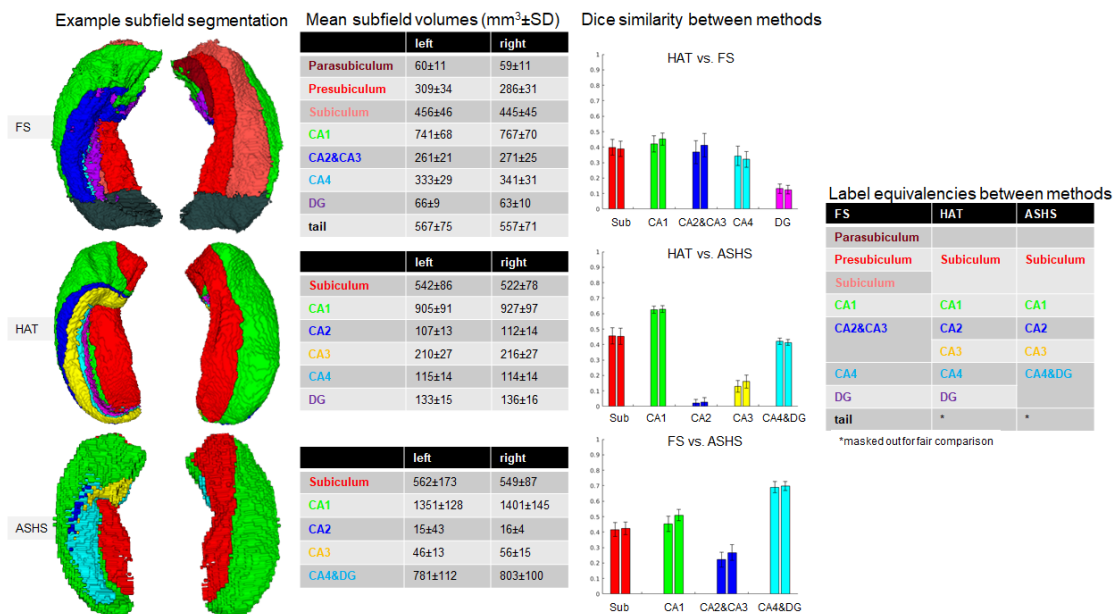


Figure 24. Subfield definitions using the method proposed here Hippocampal AutoTop (HAT), Freesurfer6.0 (FS), and Automatic Segmentation of Hippocampal Subfields (ASHS). The first column shows a single subject example (HCP1200, 105923) subfield definitions using each method. The second column shows tables of

the average subfield volumes using each method. The third column shows the Dice overlap between equivalent subfields from each method. The final column illustrates the equivalency between methods.

We examined the overlap between equivalent subfield labels between FS, HAT, and ASHS by calculating Dice overlap. The equivalence between different labels are summarized in the table to the far right of Figure 24. Briefly, we tried to match as many equivalent labels as possible between each method. However, parasubiculum and the hippocampal tail from FS had no equivalent labels and so for fair comparison, parasubiculum was discarded and the hippocampal tail was masked out of other methods when comparing to FS segmentations. Dice overlap is shown in the third column of Figure 24, and in general scores were quite low. Within the same protocol, good values might range from 0.7-0.9, or 70-90% overlap, but between these protocols no such overlap was observed for any subfield. It should also be noted that Dice overlap is generally lower for small structures, which are less likely to overlap. Thus, combining labels leads to much higher Dice overlap scores, and so the scores reported here might be considered upper bound to what the overlap of individual subfields from each method could actually be if they were present across all methods.

We further aimed to examine any systematic differences between HAT and the other two methods (FS and ASHS) in more detail. For this, we quantified the 3D deformations of label-label registrations between HAT and FS, and HAT and ASHS (using the same label equivalencies as above). A single subject example of such a deformation is shown in Figure 25A, including the same coronal and sagittal slices segmented using each of the three methods. Using unfolded space, we were able to average these deformations across subjects despite gross differences in shape in native space. We show these averaged deformations in Figure 25B as vectors, the lengths of which correspond to real-world distances. All subfields of the hippocampal head extended more posteriorly in FS compared to HAT, and the body of the hippocampal was generally wider. That is, subiculum extended more medially in FS compared to HAT, while the lateral subfields CA1 and CA2&CA3 extended even further laterally. Deformations were less extreme

when comparing ASHS to HAT, but all subfields of the head and tail were generally found to extend further posterior.




Figure 25. Deformations between equivalent subfield produced using the method proposed here Hippocampal AutoTop (HAT), Freesurfer6.0 (FS), and Automatic Segmentation of Hippocampal Subfields (ASHS). A) shows equivalent labels from each method from a single subject, as well as the fluid label-label 3D deformations between them. Shown here is the example of the left hippocampus of HCP1200 subject 105923. B) shows the group-averaged deformations from a HAT to each other method (right), as well as the mode subfield identities projected into unfolded space and on a group-averaged mid-surface (left).

One possible contributor to differences in the posterior hippocampal head, sometimes referred to as the apex of the uncus, may have to do with partial voluming with surrounding CSF, blood vessels, and other white matter structures. This can be seen in the sagittal slices of the example case shown in Figure 25A, where FS in particular includes some of these more posterior structures as part of the hippocampal head. In the tail of the hippocampus, small and frequent digitations are often seen (again visible in

sagittal slices of Figure 25A). This is not addressed in FS which does not distinguish subfields in this region, and in ASHS these smaller scale digitations are not segmented separately. In the example shown, these small digitations were not all delineated in HAT either, and indeed the SRLM that separates these digitations may have been overestimated (and therefore excluded from the analyses shown here) due to partial voluming.

Another key difference between the protocols examined here is that in the anterior hippocampus, HAT shows subiculum wrapping around the antero-medial hippocampal head and running along the superior aspect of the hippocampus to the vertical component of the uncus, while in FS and ASHS subiculum is strictly confined to the inferior aspect of the hippocampus. It was recently noted by (Ding and Van Hoesen, 2015) that the subiculum is consistently present on the superior aspect of the hippocampus, even across variable folding patterns in the hippocampal head.

Finally, we can examine the subfields generated under each protocol relative to the topology (or unfolded space) as estimated through HAT (Figure 25B, left). HAT shows clear distinction between all subfields, but ASHS showed no CA2 and little CA3 throughout most of the anterior-posterior extent of the hippocampus. This is because the subfields were not consistently aligned to the estimated folding of the hippocampus, and so the mode label at each unfolded point was not aligned in these smaller subfields. For the same reason, we also see discontinuities in labels CA2&CA3, CA4, and DG under the FS protocol. Only HAT showed clearly defined and contiguous subfields along the topological folding of the hippocampus. However, it should be noted that HAT inherently relies on this topologically unfolded space for defining subfields and so it is to be expected that this method alone respects the known topological contiguity of the hippocampal subfields.

4.4 Discussion

4.4.1 Methodological approach

The goal of the methods explored here was to combine the power and flexibility of deep learning based segmentation with the constraints, flexibility, and feature extraction utility

of a surface-based segmentation approach. Deep learning has come to dominate image processing in many domains, and especially U-Net or similar variants in medical image segmentation. However, U-Net is often limited by the availability and consistency of data and manual ground truth segmentations. The present study pooled data across several initial studies employing our surface-based unfolding (Adler *et al.*, 2018; DeKraker *et al.*, 2018, 2020; Ferko *et al.*, 2017), and we employed incremental learning to gradually achieve good ground truth segmentations for the full HCP1200 dataset. This yielded broad hippocampal tissue classes (i.e. grey matter, SRLM, and surrounding structures but not hippocampal subfields) that were as good or better than manual segmentation in this dataset (Figure 20). That is, Dice overlap between U-Net and manual counterparts in a left out test set were higher than what is typically seen between two manual inter-raters. This high level of performance remained true even after post-processing via template shape injection which was required to enforce the conditions for subsequent topological unfolding. We explored the consequences of parameter choices (template selection and update field variance) of this post-processing in Figure 21, which justified our choice of parameters and illustrated the robustness of template shape injection in this context.

We next applied our previously developed hippocampal unfolding framework and feature extraction. This approach was motivated by the 2D archicortical organization of the hippocampal subfields. By delineating the folding of this mantle, hippocampal tissue can be indexed and flatmapped, constraining problems like subfield segmentation from 3D to 2D. This is a clear analogue to state-of-the-art methods for neocortical parcellation (see Dale, Fischl and Sereno, 1999; Fischl, Sereno and Dale, 1999; MacDonald *et al.*, 2000; Zijdenbos, Forghani and Evans, 2002; Kim *et al.*, 2005; Glasser *et al.*, 2016; Research and Case Medical Research, 2019). The exact definitions of subfields in 2D are still widely debated, but as a proof of concept we applied the same definitions to all subjects, which were motivated by previous 3D histology work (DeKraker *et al.*, 2020). These subfield definitions are then projected back to each subject's native space, which can vary considerably depending on that subject's pattern of hippocampal folding. An average of this folding as well as a single subject example are shown in Figures 4 and 5, respectively. Features extracted under this unfolding protocol, including thickness,

gyrification index, and T2w, differed markedly between the different subfields, in agreement with previous work (DeKraker *et al.*, 2018, 2020) (Figure 22).

4.4.2 Comparison to extant methods

We benchmarked the subfield segmentation proposed here against other popular methods Freesurfer6.0 (FS) and Automatic Segmentation of hippocampal Subfields (ASHS) to highlight systematic differences in the results obtained using these methods. Note that there is no direct validation dataset to which we can compare any of these results, and so instead we evaluate each method by comparison to extant anatomical literature. Each of these methods use a different labelling scheme, and so we drew direct comparisons only where anatomically appropriate (see Figure 24 right). Broad discrepancies in Dice overlap were seen between all three methods, which highlights the need for harmonization of methods (Yushkevich, Amaral, *et al.*, 2015; Wisse *et al.*, 2017) and also suggests that caution should be exercised when comparing subfields across these methods. In general, total volumes were similar (Figure 23) but each subfield volume was smaller under our proposed Hippocampal Automated Topology (HAT) method than under FS or ASHS (Figures and 6). This is most likely attributable to our extensive delineation of SRLM across the full anterior-posterior extent of the hippocampus, which was limited in FS and not present at all in ASHS. Partial voluming with surrounding structures and between small digitations in the hippocampal tail may also contribute to the overall volume discrepancies of subfields between methods. As discussed in previous work (DeKraker *et al.*, 2018, 2020; Cai *et al.*, 2019), failing to delineate digitations could lead to overestimations of hippocampal thickness and volume, with a corresponding underestimation of SRLM. This was likely the case in FS and ASHS, and to some extent this was true in HAT as well, particularly in the small digitations often present in the hippocampal tail (e.g. see comparison to DeKraker *et al.*, 2020).

Some systematic differences between FS, ASHS, and HAT can also be seen in the transverse, or unfolding, directions. Notably, the subiculum-CA1 boundary varies considerably between protocols, which can be seen in coronal slices or in the unfolded space in Figure 25. That is, this boundary is generally more medial under ASHS and more lateral under FS, with the HAT definition being in between. Part of this broad

discrepancy may have to do with the distinction of prosubiculum as being a distinct subfield and part of the subicular complex, or simply a transition area to be divided between subiculum and CA1 (Wisse *et al.*, 2017). Other discrepancies in the transverse direction were less extreme, except in the anterior and posterior edges of the hippocampus which were labelled predominantly CA1 under both ASHS and FS (and no label is provided in the posterior tail of the hippocampus from FS). This may be an oversimplification on behalf of ASHS and FS since recent and long-standing work (e.g. Duvernoy, Cattin and Risold, 2013; Ding and Van Hoesen, 2015) shows that all subfields are present even in the most anterior vertical component of the uncus. Finally, it should be noted that only HAT showed consistent subfield contiguity in the unfolded or transverse direction across the full anterior-posterior extent of the hippocampus. This was necessarily the case since HAT subfields are contiguous defined in unfolded space, but this contiguity is well described in histology and can readily be enforced under our protocol.

4.4.3 Limitations and future directions

As discussed above, one limitation of the current approach is that there remain small-scale digitations in hippocampal grey matter that are not always captured using the current method (notably in the hippocampal tail, but some digitations in the remainder of the hippocampus may also be attenuated). This likely stems from two methodological limitations: firstly, resolution limits the contrast of the SRLM and alveus which separate digitations, particularly at a small scale. Secondly, the post-processing step employed here entails some regression towards an averaged, smooth hippocampal morphology. This post-processing was necessary to enforce separation between digitations and across the SRLM precisely because of segmentation errors that are most common in ambiguous parts of the hippocampus, like the SRLM in the hippocampal tail. Future improvements could include more detailed segmentation data for U-Net training, which could entail more careful manual segmentation of the SRLM in the hippocampal tail, or improvements in acquisition resolution. Further U-Net training could also alleviate the need for the post-processing employed here, eliminating the small amount of regression towards a smoother, averaged hippocampal shape. Note that details like small scale

digitations and corresponding shifts in subfield boundaries are not captured by other manual or automated protocols for hippocampal subfield delineation. Thus, we believe that the overall pipeline proposed here is the most detailed protocol for *in-vivo* hippocampal analysis to date.

In the current work, we applied the same hippocampal subfield boundaries to all subjects in unfolded space. A more subject-specific approach may vary these boundaries in 2D according to further available information, such as the quantitative and morphological features extracted and plotted in unfolded space. In native space this would constitute a topological shift in subfield boundaries. This may help account for further subject-specific variability, but could also introduce systematic errors, especially if measures like thickness are inflated due to imperfect delineation of digitations or the presence of abnormal tissue as in some diseases differentially affecting hippocampal subfields (e.g. Blümcke *et al.*, 2013; Steve *et al.*, 2020).

Generalizability was high between the training and test sets examined here, but it remains unclear how HAT will perform on different datasets with different resolutions and contrasts. In particular, the convolutional kernels applied in U-Net are well suited for the current data but may not be appropriate for higher resolution images with sharper tissue boundaries, lower resolution images with smoother tissue boundaries or other contrast types (e.g. T1w where contrast is inverted). Other deep learning applications use fine-tuning to overcome analogous issues in generalizability, wherein a pretrained network is further trained on a new target dataset. This can be combined with an incremental learning approach to limit the requirement for manually segmented training data. That is, the current HAT pipeline can be applied to a new dataset, and then U-Net can be further trained using only the subset of the new dataset that was successful as assessed by manual inspection. This subset could be expanded further by manually adjusting U-Net outputs that were close but not perfectly segmented. This approach has been highly successful in other deep learning applications with limited data, as in a research context, such as biological motion tracking (Nath *et al.*, 2018; Mathis *et al.*, 2018; Graving *et al.*, 2019; Heras *et al.*, 2019). However, it should be noted that the success of the current methods and the level of detail therein will depend critically on the visibility of folds,

such as hippocampal digitations and the SRLM, in new datasets which likely requires sub millimetric resolution (DeKraker *et al.*, 2020).

The current code, including pre trained U-Net model, are all available for deployment in new research contexts via a containerized BIDSapp or a MATLAB-based development version at https://github.com/jordandekraker/Hippocampal_AutoTop. This code is modular and well documented for easy expansion to new domains, and future releases will include U-Net models trained on new datasets with different resolution and contrast properties. Some automated checks are employed to ensure there are no gross segmentation errors, but users are advised to manually inspect results to ensure good quality on new datasets. Tools to easily perform such visualizations are also included. Application to new datasets is advantageous for answering questions about those data as well as for expanding the generalizability of the methods presented here. Thus, it is our hope that in future work the generalizability, detail, and feature extraction capability of this hippocampal automated topology framework will only grow, especially with increased availability of high resolution MRI data or even 3D data from other acquisition methods.

4.4.4 Conclusions

In the current work, we present a method termed Hippocampal Automated Topology which performs tissue segmentation, unfolding, feature extraction, and subfield definition of sub millimetric T2w MRI data. This method combines the power and flexibility of deep learning with subsequent constraints and utility afforded by surface-based approaches to cortical segmentation. We developed this method specifically for 0.7mm^3 T2w data found in the open source Human Connectome Project (HCP1200) dataset. This method produces segmentations that are as good or better than manual segmentations in terms of Dice overlap, and all HCP1200 cases were amenable to subsequent unfolding and feature extraction. In agreement with previous work, extracted features including thickness, gyrification index, and T2w differed markedly between subfields. Comparison to other popular automated hippocampal subfield segmentation methods reveal systematic differences which, for the most part, favour the methods proposed here. This method can readily be extended to new datasets with different resolution and contrast.

We hope that with automation of these methods, more labs will consider the highly variable folding in the hippocampus which is critical for segmenting or extracting features from hippocampal substructures (e.g. subfields, individual digitations, or other subdivisions like laminae or anterior-posterior differentiation) in detail.

Chapter 5

5 General Discussion

5.1 Background and scope

The research conducted in this thesis represents an effort to understand the true underlying folding structure of hippocampal tissue in neuroimaging. This research was initially motivated by an attempt to mechanistically link functions of the hippocampus to microcircuits that are stereotypical of the individual hippocampal subfields (see Chapter 1.2.2). In other words, the initial goal of my graduate work (and the current direction of the field of hippocampal research more broadly) was to decompose the overall functions of the hippocampus at the macroscale into sets of microscale elements. As discussed in Chapter 1.4.3, this can be done on a rudimentary level by overlaying subfield definitions from *ex-vivo* histology onto *in-vivo* MRI images containing much less detail. However, in this endeavour it quickly became apparent that mesoscale differences at the level of folding and gyrification between individuals increases the complexity of this problem. All three projects of this thesis are thus devoted specifically to this problem, as addressed using a topological coordinate framework.

The problem of mesoscopic folding within the hippocampus is not ubiquitously recognized in the field of hippocampal neuroimaging. As the promise and feasibility of hippocampal subfield delineation in neuroimaging became clear, many manual and automated protocols were developed in quick succession to estimate subfield boundaries (Yushkevich *et al.*, 2015). Initially, none of these protocols placed emphasis on the folding of the hippocampus, which was assumed to remain relatively constant in coronal slices through the hippocampal body. Folding in the hippocampal head and tail was frequently simplified or excluded in these protocols because it was deemed too complex and beyond the capabilities of these protocols. Nevertheless, many differences between protocols were clear even within only the body of the hippocampus, which have given rise to concerns about replicability and comparability between studies performed using disparate protocols. To address these concerns, an international effort to harmonize the protocols has been launched (Wisse *et al.*, 2017; Olsen *et al.*, 2019). Concurrently, and

with these concerns in mind, new investigations using *ex-vivo* histology and MRI have begun to elucidate the complexity and inter-individual variability seen in the arrangement of hippocampal subfields. This variability is currently receiving increased attention among researchers, with likely links to aging (Cai *et al.*, 2019) and disease (Oppenheim *et al.*, 1998; Henry *et al.*, 2011). The methods described in this thesis were developed in an effort to address these complexities and inter-individual differences under a common 2D surface-based framework.

5.2 Summary of projects

In Chapter 2, we pursued a computational approach to take advantage of regularities in hippocampal structure related to ontogeny, which entails consideration of the hippocampus as a folded archicortical structure. Through computational unfolding of the hippocampus, the proposed protocol provides a coordinate system that can index hippocampal tissue in a precise and flexible manner, while capturing the noticeable inter-individual differences in morphology that have been documented in histological studies of this structure. This method critically depends on the visualization of the SRLM, or ‘crease’, along which the hippocampus is folded. We argue that this method offers several practical advantages over manual segmentation techniques. These advantages can be summarized as follows: i) Unfolding hippocampal grey matter allows for indexing of analogous tissues (or sets of candidate boundary locations) across participants with variable morphologies. ii) The unfolded coordinate space can be used for inter-subject alignment and subsequent mapping of properties across the full long-axis and proximal-distal extent of the hippocampus, as illustrated in Chapter 2.3.3 for intracortical myelin and cortical thickness measures. iii) Segmentations applied in this unfolded coordinate space show good spatial overlap with, and may even correct for tracing errors in, detailed manual subfield segmentations. This coordinate system also captures subtle but critical structural features, as demonstrated in a direct comparison with a resected histological sample from a patient with epilepsy (see Chapter 2.3.4).

In Chapter 3 we mapped the human hippocampus in detail by combining three methods. First, we used a unique dataset, BigBrain, that contains both histological-level detail and fully 3D spatial context. Second, we imposed the topological hippocampal unfolding

framework developed in Chapter 2. Third, with this framework we extracted a set of morphological and laminar features, the latter of which have been used prolifically in neocortical characterization and parcellation. Using these methods, we made three novel empirical observations: i) Unsupervised clustering of these features closely resembled classically defined hippocampal subfields. ii) Despite traditional reliance on laminar features in histology, morphological features alone were sufficient to closely approximate most hippocampal subfields. iii) Some features such as gyrification in CA1 showed, at least qualitatively, subfield-specific anterior-posterior differences that might relate to functional differences described in the extant literature. Overall, these findings highlight new structural characteristics of the hippocampus, and offer promising avenues for improved delineation and characterization of hippocampal subfields using *in-vivo* neuroimaging.

In Chapter 4, we presented a method termed Hippocampal Automated Topology (HAT) which performs fully automated tissue segmentation, unfolding, feature extraction, and subfield definition given sub millimetric T2w MRI data. This method combines the power and flexibility of deep learning with subsequent constraints and utility afforded by the unfolded coordinate framework developed in Chapter 2. We developed this method specifically for 0.7mm³ T2w data found in the open source Human Connectome Project (HCP1200) dataset. This method produced tissue segmentations that showed high Dice overlap with manual segmentations, and all HCP1200 cases were amenable to subsequent unfolding and feature extraction. In agreement with previous work, extracted features including thickness, gyrification index, and T2w differed markedly between subfields. Comparison to other popular automated hippocampal subfield segmentation methods revealed systematic differences which, for the most part, favored the methods proposed here in terms of detail and qualitative comparison to anatomical literature. This method can readily be applied to new sub millimetric T2w datasets, and supporting methods are offered for generalization to new datasets with different resolution or contrast. We hope that with automation of these methods, more researchers will consider the highly variable folding in the hippocampus which is critical for segmenting or extracting detailed features from hippocampal substructures like subfields, individual digitations, laminae, or anterior-posterior gradients.

5.3 Accommodation of specific hippocampal features in the unfolding framework

5.3.1 Curvature in the hippocampal head and tail

As discussed in Chapters 1 and 2, we aimed to computationally fit the gross curvature seen in the hippocampal head, including the uncus and vertical component of the uncus, as well as smoother curvature in the hippocampal body and tail. With this gross curvature accounted for, we found that similar subfield definitions could be applied across the full anterior-posterior extent of the hippocampus, which is most easily visualized in unfolded space. This aim was consistently met throughout all applications of our unfolded coordinate framework in this thesis which is briefly illustrated in Figure 26. However, investigations at higher resolution (ex-vivo MRI and 3D histology) revealed several additional smaller-scale folds or digitations that were not well captured in manually segmented or subsequently automatically segmented (Chapter 4) MRI images. This is discussed in the following section.




Figure 26. Mid-surface models of hippocampal folding estimated through MRI (A) and 3D histology (B). Purple arrows indicate folds seen in histology that are not typically seen using *in-vivo* MRI, while black arrows represent gross curvature and digitations that are typically seen in both modalities. Digitations in the hippocampal body are not always present and are absent from this MRI sample. Such digitations

are typically visible but likely attenuated in MRI under current acquisition protocols.

5.3.2 Fine scale folding in the uncus and tail

Long standing anatomical work has illustrated a distinct fold within the vertical component of the uncus. This fold mimics the classic ‘C’ shape of the hippocampal subfields in the remainder of the hippocampus but at a smaller scale and in the axial, rather than coronal, plane (Duvernoy, Cattin and Risold, 2013). This fold was not well captured in our manual segmentations and subsequent unfolding protocol using *in-vivo* MRI, as illustrated in Figure 26. This limitation was primarily due to limitations in image resolution and contrast. While high myelin SRLM does exist to differentiate this fold, it was not consistently visible in the dataset examined. In Chapter 3, resolution was well above the level of individual folds within the hippocampus, and we have high confidence that the full topology of the hippocampus was correctly delineated and unfolded, even in the vertical component of the uncus. In Chapters 2 and 3 we noted that digitations were not confined to the hippocampal head as previously described (Ding and Van Hoesen, 2015), but rather in many cases extended throughout the length of the hippocampal body, with some other recent studies making similar observations (Cai *et al.*, 2019; de Flores *et al.*, 2020). Indeed, with the microscopic levels of resolution used in Chapter 3, these digitations could be seen extending throughout the hippocampal tail, where they were much smaller than the more anterior digitations (see Figure 26). Reexamining *in-vivo* data from Chapter 2 while keeping in mind the digitations noted in Chapter 3, it is still not clear where the fold within the vertical component of the uncus is located, but it is clear in some cases that smaller scale digitations are present within the hippocampal tail.

The overall effect of failing to delineate digitations or folds, whether in the vertical component of the uncus or in the body or tail of the hippocampus, is an overestimation of thickness and an underestimation of surface area. This underestimation of surface area also causes topological shifting of the hippocampal subfields in unfolded space. In particular, we noted in Chapter 3 that digitations tended to be confined mostly to CA1. Thus, if we apply boundaries from a sample with an extensively folded CA1 to a sample which showed no delineated digitations in CA1, we will overestimate the true extent of

CA1, and all other subfields will be slightly shifted, proportionally. In the case of the fold in the vertical component of the uncus, a similar problem is present: by failing to delineate this fold the proximal distal axis in this region essentially forms a line instead of a hook or flattened ‘C’ shape. As with the hippocampal tail, subfields defined on a more detailed map of hippocampal folding will become topologically shifted when applied to a simpler, less folded map of hippocampal folds.

A U-Net deep learning architecture was used to automate tissue segmentations in Chapter 3 (including alveus and SRLM that separate the folds and digitations of the hippocampus). U-Net was trained with manual segmentations from the same protocol and some of the same data as Chapter 2. Furthermore, the HCP1200 dataset was slightly lower in resolution than the data used in Chapter 2, and so it is unlikely that every digitation was delineated. Finally, we also included a post-processing step which entails some regression towards a smooth, average hippocampal shape and therefore subtle reduction in the magnitude of digitations. This step was necessary to avoid major distortions of the coordinate framework due to imperfect U-Net segmentation by removing topological discontinuities or erroneous ‘bridges’ between tissues. Thus, while perfect segmentation of all mesoscopic folds in the hippocampus is close, and may be achievable in the near future, the methods presented in Chapters 2 and 4 here still show some room for improvement. This will most critically depend on resolution which limits our ability to detect small scale folding, most notably in the vertical component of the uncus and the hippocampal tail. This improvement could be accomplished through fine-tuning the U-Net architecture employed in Chapter 4, which I will discuss further in Chapter 5.4.

5.3.3 Non-topologically defined hippocampal substructures

The intra-hippocampal structures SRLM and DG are not unfolded in our current framework, mainly due to limitations in image contrast and resolution. Instead, these distinct structural labels are only employed to demarcate the boundaries of hippocampal grey matter. That is, SRLM marks the inner edge of hippocampal grey matter (continuous with the pial surface of the neocortex) while the DG marks the distal edge of hippocampal grey matter. SRLM is typically considered laminae of the CA fields (with

SRLM extending as laminae of the subiculum) and together this structure is composed of dense axonal and dendritic fields high in myelin and blood vessels. These properties give rise to differing contrast in MRI, bearing greater resemblance to white matter than grey matter. Indeed, this contrast difference is the criterion for manually or automatically defining SRLM in most protocols. Our current framework could be extended to map out the SRLM via a process of interpolating its nearest corresponding points in hippocampal grey matter, as in (Hussain *et al.*, in preparation), but without microscopic resolution it is difficult to determine which parts of SRLM belong to which surrounding grey matter folds (for example, a given SRLM voxel in the hippocampal body may be a part of CA1 which is lateral to it, or CA4 medially to it). Though it presents challenges, this may be a useful future direction, especially given that some work has suggested that the quantitative properties of SRLM may be a useful biomarker in Alzheimer's disease (Kerchner *et al.*, 2010).

No contrast differences were seen between the DG and its neighbouring CA4 in MRI, and thus under our manual protocol this boundary had to be defined heuristically. This was challenging since the DG is composed of a thin granule cell layer with its own topology, distinct from the rest of the hippocampus. However, the DG does closely follow the SRLM including between the digitations where SRLM folds outward into digitations (see Duvernoy, Cattin and Risold, 2013). Thus, the heuristic we used to define the DG was a dilation of the SRLM over the distal portions of hippocampal grey matter by a distance of 0.3mm. This roughly corresponds to the thickness of the healthy DG, but this structure is itself complex as it occasionally shows additional inward folding upon itself (Duvernoy, Cattin and Risold, 2013). These same rules were used in the manual segmentation on which U-Net was trained in Chapter 4. Thus, under the methods employed in Chapters 2 and 4, the DG will consistently show a thickness of approximately 0.3mm but may vary in width depending on the morphology of the SRLM. Since our protocol involves detailed delineation of the SRLM, and the DG is tightly anchored to SRLM, we assert that our method still holds high validity compared to other currently utilized segmentation protocols, but it is important to note these limitations, particularly in estimating the thickness and folding within DG. Finally, it should be noted that some *ex-vivo* work has found distinct MRI intensities between the DG and CA4.

Thus, it may be possible to define this structure in a data-driven rather than heuristic fashion if future improvements in MRI contrast and resolution are made.

5.4 Adaptation of current approach to new data

5.4.1 Higher resolution datasets

MRI has steadily improved in resolution and contrast since its integration into medical imaging and basic research, and with several promising lines of investigation into new acquisition methods, this trend seems likely to continue. This may obviate certain hippocampal features, such as the SRLM that separates the different folds of the hippocampus. Thus, technological improvements might highlight the need to consider inter-individual differences in intra-hippocampal folding in future neuroimaging work focused on the hippocampus. This is not typically explicitly considered in current methods such as Freesurfer (FS), Automatic Segmentation of Hippocampal Subfields (ASHS), or the host of manual subfield segmentation protocols developed in recent years. However, as noted above, the methods presented here still do not perfectly delineate all folds or digitations within the hippocampus. As illustrated in Chapter 3, the general coordinate framework we have developed can indeed be applied to a fully topologically correct hippocampal tissue segmentation if available. Therefore, the main element that could be improved with higher resolution imaging is the manual or automated U-Net segmentation of hippocampal tissue classes. In anticipation of this eventuality, and to better fit images of different contrast types, we have included recommendations and code for how to go about fine-tuning U-Net to fit new datasets or new manually defined structural features as they become available. Briefly, this could be done using incremental learning which would entail applying the current pipeline and then manually inspecting and correcting results until good quality tissue segmentations are achieved for a subset of the new dataset. The U-Net architecture can then be trained on this new subset of data and then tested on the remainder. Ideally this would result in good segmentation for the entire novel dataset, though it remains unclear how large this dataset should be and how much manual or semi-manual correction would be required. This will most likely depend on the extent to which new datasets differ from the ones on which U-Net is currently trained, in resolution and in contrast.

5.4.2 Anisotropic datasets

As discussed in Chapter 1.4.2, MRI for hippocampal subfield segmentation is often performed using anisotropic or thick slice T2w scanning, with high in-plane coronal resolution (e.g. $0.4 \times 0.4 \times 2.0 \text{mm}^3$). This makes it easier to view the SRLM in coronal slices through the hippocampal body, but more difficult in the hippocampal head and tail and furthermore this may obscure digitations. Nevertheless, good contrast is seen in many such acquisitions, and it may still be possible to discern or infer out-of-plane hippocampal features such as its curvature and digitations in these data. The current proposed methods involve resampling all images to 0.3mm^3 , which generates significant blurring of anisotropic data in the anterior-posterior direction. As with datasets using different resolutions, discussed above, the current U-Net segmentation method is not optimized for this level of contrast, and could likely be improved through further training or fine-tuning on such data.

5.4.3 Datasets with T1w or other contrasts

As discussed above, the unfolding framework presented in this thesis can be applied across different resolutions and image modalities, provided a detailed segmentation of the different hippocampal tissue types (grey matter, white matter alveus and SRLM, and surrounding structures) can be attained. Indeed, the current thesis applied the same methods to *in-vivo* T2w MRI, *ex-vivo* T2w MRI, and *ex-vivo* 3D histology. The current automated application, HAT, is optimized specifically for T2w, but can flexibly be retrained or fine-tuned to suit other contrast types.

Though T2w MRI is most common in hippocampal subfield delineation, T1w structural imaging is more common for imaging the rest of the brain in part because of the high contrast it provides between neocortical grey matter and underlying white matter. Thus, in studies where both the hippocampus and also other brain structures are of interest, T1w may be the only or the best resolution and contrast imaging modality available, and some subfield segmentation protocols can be applied to T1w images alone (e.g. Iglesias *et al.*, 2015). In particular, T2w imaging shows good contrast in the hippocampal SRLM while in T1w imaging this feature exists but is not as clear (see Figure 2 for example). As

discussed in Chapter 1.4.1, there is an imprecisely inverse relationship between T1w and T2w. That is, these two modalities are typically inversely related, but they also show differences in their susceptibility to molecules other than fat and water, such as iron sequestered inside grey matter tissue or blood. If not for these tissue contrast differences, it would be relatively straightforward to apply the current version of Hippocampal AutoTop to T1w data by simply computing inverse T1w. Thus, it is most likely that fine-tuning or retraining of the U-Net tissue segmentation model would be required before application to T1w images, as overviewed in Chapter 5.3.1. Further improvements in T1w contrast and resolution may make this modality even more desirable, especially if additional scans are not required for imaging neocortex and hippocampus separately. The additional time that this affords could be spent increasing resolution, or acquiring additional T1w images for averaging, improving contrast as discussed in Chapter 1.4.2.




Figure 27. Sample scans employing T2w, T1w, and low b-value dMRI acquisition protocols. The T2w scan represents a different subject, since no directly comparable data were available. Note that CSF and white matter within and surrounding the hippocampus show opposite intensities in T2w and T1w images, but not in dMRI images.

5.4.4 Diffusion magnetic resonance imaging

Another MRI acquisition technique, diffusion MRI (dMRI) shows particularly high sensitivity to structural abnormalities in the hippocampus (Förster *et al.*, 2012). dMRI measures the diffusion of water, which is constrained by structures like cell membranes.

Thus, anisotropy in diffusion is seen in structures like heavily myelinated axon bundles, such that diffusion is greatest in parallel to the direction of these bundles. The resulting anisotropies can be used to estimate the course of large fiber bundles, a technique called tractography. Such tracts have been successfully traced within and around the hippocampus in recent work in both health (Augustinack *et al.*, 2010; Yassa, Muftuler and Stark, 2010; Adnan *et al.*, 2016) and disease (Pereira *et al.*, 2014; Dinkelacker *et al.*, 2015). This work has improved our understanding of hippocampal connectivity and revealed degradation of some pathways in aging and disease. Diffusion properties can also differ between grey matter tissues with different properties, like the density of dendrite networks which also act to constrain diffusion. This is the principle behind intracortical quantitative dMRI, which can be used to distinguish neocortical regions with stereotyped cytoarchitectonic properties (Aggarwal *et al.*, 2015; Ganepola *et al.*, 2018) and even to distinguish laminae within the neocortex (Leuze *et al.*, 2014). In particular, this can be done by differentiating diffusion along the columnar or transverse direction from diffusion along the laminar or inner-outer direction. One prerequisite of such an approach is having a good model of the direction of cortical folding, and so in recent and ongoing work we have developed a method to apply our hippocampal unfolding framework to dMRI data (Hussain *et al.*, in preparation). This involves mapping not only the different positions within the hippocampus to an unfolded space, but also rotating corresponding directional encoding data. Briefly, this work used high quality *in-vivo* dMRI acquisitions and some key findings showed that unfolding the hippocampus leads to a more anatomically plausible modelling of the connectivity of the hippocampus as probed by probabilistic tractography, revealing key elements of the polysynaptic pathway and anterior-posterior connectivity gradients. Future work along this line of investigation will further examine intracortical measures along the columnar and laminar directions within the hippocampus.

Finally, dMRI is typically performed at lower resolution (e.g. $>2\text{mm}^3$) than other structural scans (e.g. $<1\text{mm}^3$), but recent work has developed promising acquisitions for high resolution dMRI (1mm^3) that may even be useful in detecting hippocampal folds or digitations (Treit *et al.*, 2018). This may be useful because it overcomes an important limitation in T1w and T2w imaging: cysts or pockets of CSF are often seen along the

hippocampal sulcus, between laminae of the SRLM. In T1w and T2w imaging, this gives rise to partial voluming of CSF with white matter, which can thus resemble grey matter. This makes the SRLM and, subsequently, the folds of the hippocampus difficult to detect unless very high resolution is achieved. Under the protocol proposed by (Treit *et al.*, 2018), higher resolution is achieved through low b-value (or less directional encoding), and segmentation is performed on the mean diffusion image rather than in any given directionally coded volume. dMRI under this novel protocol shows similar contrast between white matter, as in the SRLM, and cysts and so while it remains difficult to distinguish cysts from SRLM, it is relatively straightforward to distinguish both of these structures from hippocampal grey matter and therefore understand intra-hippocampal folding. Current and future work could aim to further investigate similar acquisition techniques and their utility for our hippocampal coordinate framework (for example, see Figure 27).

5.5 Applications

5.5.1 Research in patient populations

Having more accurate and detailed delineation of hippocampal subfields can improve the quality of hippocampal neuroimaging in many respects, including quantitative MRI, morphometry or volumetry, and fMRI. Extensions of the work described here have already begun on neuroimaging patients with Major Depressive Disorder (MDD) and, separately, patients with medial temporal lobe epilepsy (mTLE). Both of these studies were designed to make use of the hippocampal unfolding and analysis methods presented here (that is, they include high resolution isotropic T2w MRI scans). Different subtypes of mTLE have been shown to differentially affect hippocampal subfields CA1 and CA4 in terms of neuronal loss and overall volume or atrophy (Blümcke *et al.*, 2013). More recent work has suggested these mTLE subtypes may even be detectable using *in-vivo* imaging (Blumcke *et al.*, 2017), which could prove useful in treatment planning. The current work improves subfield definitions over previous work and extends these definitions into the hippocampal head and tail, which was specifically identified as a shortcoming of previous mTLE subfield imaging (Blumcke *et al.*, 2017). Volume losses specific to each subfield can be detected in this way, and these measures can also be

further decomposed into surface area and thickness, one or both of which may be more sensitive than overall volume measures of structural differences between patients and healthy controls (or between the subtypes of mTLE patients). Neuronal loss is difficult to detect using MRI; however, this can sometimes entail the formation of sclerotic (or scar-like) tissue within hippocampal grey matter. This can be detected in quantitative MRI as a change in intracortical myelin or diffusivity compared to healthy tissue, as discussed above. These measures are ameliorated under our unfolding framework since MRI values can be sampled along a mid-surface, reducing the partial voluming of other nearby CSF or white matter structures. This is a critical improvement over currently implemented measures since the contrast between grey matter and CSF or grey matter and white matter is greater than that expected between sclerotic tissue and healthy tissue, especially if sclerosis is in early stages (Jackson *et al.*, 1990; Berkovic *et al.*, 1991; Cross *et al.*, 1993; Coan *et al.*, 2014). Accidental inclusion of these surrounding structures can drastically bias results, adding noise to experimental analyses or, in cases of systematic morphological differences between the populations being compared, this could lead to systematic errors. Thus, in applying our hippocampal unfolding framework to MRI data obtained from patients and healthy controls, we expect to see improved sensitivity compared to previous work and can concurrently examine any systematic differences in morphology to identify potential systematic biases.

Similar principles apply to hippocampal imaging of patients with MDD. In this case we are particularly interested in the DG subfield as the primary site of adult neurogenesis. Some animal and human imaging studies have already shown links between MDD and hippocampal subfields, of which the DG is most notable (Maruszak and Thuret, 2014; Malykhin and Coupland, 2015). Some of these links even include increases in volume with recovery from depression, which could be in part driven by neurogenesis, though contributions from other processes like synaptogenesis or angiogenesis are likely (Firth *et al.*, 2018; Sahay and Hen, 2007). The ongoing study employing the methods presented here aims to collect behavioural data on tasks specifically targeted at hippocampal subfields functions, like pattern separation in the DG as discussed in Chapter 1.3.1. Thus, improved delineation of the DG may also help elucidate its functional contributions to behavioural measures. Specifically, improved DG measures improve the sensitivity of

possible links between this structure's integrity and its function as measured behaviourally. Additionally, this study employs fMRI which is often used to probe the function of the hippocampus and hippocampal subfields. Such endeavours stand to improve with better hippocampal subfield delineation in the same way that quantitative measures stand to improve through improved sensitivity and improved localization which I will discuss further below.

5.5.2 Localization beyond subfields

Another benefit of the hippocampal unfolding method described here is that it allows for localization of tissue properties beyond subfields or coarse anterior-posterior subdivisions that are typically employed. Projecting quantitative MRI or fMRI data to our hippocampal unfolded space provides implicit registration between subjects not seen in any other hippocampal analysis method. This allows for direct comparison of each unfolded point across subjects, rather than having to average values within an entire ROI. In this way, structural abnormalities can be more precisely localized. For example, a systematic analysis of all unfolded points might reveal statistically reliable effects only in proximal CA1 rather than CA1 as a whole. This precision is also afforded in the antero-posterior axis of the hippocampus. Equivalently, these unfolded pixels can be treated as vertices in a 3D surface mesh, as in the mid-surface representations shown here. This type of representation is popular in neocortical analyses, allowing for vertex-wise analysis rather than voxel-wise. This is advantageous in reducing partial voluming from surrounding white matter and potentially from averaging signals in the laminar direction in situations where laminar differences are not expected. Such analyses can thus contain cleaner signals, and there are often fewer vertices than voxels. This has the advantage of reducing the number of statistical comparison in methods like General Linear Modelling, potentially striking a better balance between control of Type I and Type II error rates. Other statistical methods can also benefit from vertex-wise analysis as well, such as diffusion map embedding – a method for neocortical analysis that doesn't make hard assumptions about boundaries between regions but rather allows signals to vary in continuous but topologically described gradients (see Paquola *et al.*, in submission; de Wael *et al.*, 2020)

This type of dense comparison is also advantageous when using statistical methods that are sensitive to gradients of quantitative differences across a tissue, rather than assuming differences are confined to parcels or subfields. Indeed, in recent work such an approach yielded graded differences over the transition from the archicortex (or allocortex) of the hippocampus to the neocortex (or isocortex) surrounding it (see Figure 28B) (Paquola *et al.*, in submission). This type of analysis aligns well with previous literature which highlights distinctive subfields that differ along the proximal-distal axis of the hippocampus, but graded differences in other properties such as widespread connection with the rest of the brain across the anterior-posterior axis of the hippocampus.




Figure 28. Extension of unfolding (or surface-based) coordinate framework to surrounding medial temporal lobe neocortical structures. A) shows an extension of our proximal-distal hippocampal axis in the hippocampus (yellow) into neighbouring neocortex (orange, red, then black) according to geodesic distance. B) shows an example of a statistical analysis (based on singular value decomposition) that employs contiguous points rather than distinct parcels or ROIs. Data from (Paquola *et al.*, in submission).

5.5.3 Integration with neocortical surface representations

Surface-based methods for neocortical analysis are already widely used, and the hippocampus, or archicortex, is fully continuous with the neocortical surface. Most neocortical analysis methods cleave the hippocampus from neocortical surfaces, and instead apply volumetric analysis methods to the hippocampus similar to those used on subcortical structures (Dale, Fischl, and Sereno 1999; Desikan *et al.* 2006; B. Fischl, Liu, and Dale 2001; B. Fischl, Sereno, and Dale 1999). Many of the same advantages of surface-based methods apply in both the hippocampus and neocortex, such as facilitated registration between subjects, unfolded or partially unfolded (sometimes ‘inflated’ in the neocortex) visualizations, and sampling quantitative or functional data along a mid-surface to avoid partial voluming. Future neuroimaging pipelines could combine surface-based or vertex-wise analyses used in the neocortex with hippocampal surfaces, reducing the number of separate analyses and files being generated from these structures. One extension of the work described in this thesis stitches together hippocampal surfaces with the inner and outer surfaces from surrounding neocortical regions (Figure 28A). This project also extends the coordinate framework we employ within the hippocampus to neighbouring neocortical regions. Briefly, this was done by extending our proximal-distal coordinates to the surrounding neocortex according to geodesic distance, while anterior-posterior coordinates were linearly extrapolated (also in a geodesic fashion). Key results of this work include graded histological differences over the extended allo-isocortical axis and simultaneous connectivity differences over the extended anterior-posterior axis. This 2D organizational framework nicely summarizes many of the histological and connectivity differences seen throughout the medial temporal lobe literature (Colombo *et al.*, 1998; Poppenk *et al.*, 2013; Strange *et al.*, 2014; Zeidman and Maguire, 2016; Plachti *et al.*, 2019).

5.5.4 Impact on cognitive neuroscience

The above sections detail how the surface-based methods developed in this thesis can improve delineation of hippocampal subfields, morphometry, or statistical analyses such as vertex-wise modelling. Each of these methodological improvements will allow the field at large to improve the sensitivity and specificity of hypothesis testing, but beyond

that the current framework also opens some novel theoretical directions. For example, there are known signals which depend on the topology of the hippocampus described here. In particular, intracranial recordings show that hippocampal theta oscillations, which are the sum of many local activations, resemble travelling waves via anatomically constrained short-ranged signaling (Lubenov & Siapas, 2009). Much of this work has been carried out in rodents, which don't show hippocampal digitations as seen in humans. Thus, it is not yet clear whether such travelling waves follow digitations, which could drastically increase their distance travelled through the hippocampus. Theta oscillations are thought to have a modulatory role, and likely organize the recruitment of neurons in support of memory encoding and retrieval. Thus, the path taken by these waves could have consequences for the functional organization of the hippocampus as well. The methods described in this thesis provide a way to delineate the full topology of the hippocampus, including its digitations, making it possible to predict how theta waves might be expected to propagate across hippocampal tissue. Specifically, this method could be used to determine the geodesic (rather than absolute) distance between intracranial electrode contacts, so that the propagation of a wave between contacts can be investigated. Additionally, the way that a single wave might propagate over the entire hippocampus could take place over a large enough time and volume that it could be detectable even using non-invasive fMRI, but more work will be required to determine whether this is methodologically feasible or not.

In addition to the tri- and mono-synaptic circuits described in section 1.2.3, topologically-constrained short-range connections within and between hippocampal subfields have been observed in human *post-mortem* data (for review, see Duvernoy, Cattin and Risold, 2013). These connections are often overlooked in current *in-silico* models of human hippocampal circuitry (e.g. Norman and O'Reilly, 2003; more recent implementations in Schapiro *et al.*, 2017 and others). The inclusion of these short-ranged connections and their constraints based on the topology of the hippocampus may improve such models in future work. In particular, both the mono- and tri-synaptic pathways are known to project to the subicular complex, but it's not clear which subregions of the subicular complex are involved. Propagation of these two pathways via short-ranged connections would suggest termini on the more proximal prosubiculum and subiculum proper before subsequent

projections from these regions to the more distal presubiculum and parasubiculum, followed by projections to the adjacent entorhinal cortex - the starting and ending point of the mono- and tri-synaptic loops. Alternatively, it is certainly possible for mid-range connections within the hippocampus to bypass some of these topologically ordered regions. Future work on the connectivity of these subicular subregions would benefit from hippocampal models which account for topological constraints of folding. For example, ongoing work shows that the computational unfolding methods proposed here can improve the detail and biological validity of tractography, in part by reducing errors due to tracts that, due to limitations in resolution and signal, appear to cross hippocampal folds (Hussain *et al.*, in preparation). With a better understanding of the mid- and short-range connectivity within the hippocampus, it will be possible to update *in-silico* models improving their biological validity and potentially revealing new emergent properties.

5.6 Conclusions

Overall, the projects presented here provide multifaceted evidence for the strengths of a surface-based approach to hippocampal analysis and subfield segmentation, as developed in this thesis. In particular, these projects demonstrated that such an approach helps account for inter-individual differences in hippocampal morphology, most notably due to different configurations of hippocampal folds or digitations. By modelling these folds under our topological coordinate framework, the hippocampus can be computationally unfolded, providing a standardized space in which subjects can be aligned despite gross differences in their native 3D morphology. This is helpful for delineation of subfields and for extracting structural features of interest, many of which are also demonstrated here to differ between subfields when examined with the sensitivity afforded by our methods. Indeed, with the resolution and contrast afforded by 3D histology, features extracted using our hippocampal coordinate framework are sufficient to derive most hippocampal subfields in an unsupervised manner. These methods are now automatically deployable in new neuroimaging work, with support for extensions to new imaging modalities. It is my hope that this framework will benefit the broader hippocampal imaging community either directly through uptake of the methods proposed here, or indirectly through their impact on the development of related new methods.

References

- Abadi, M. *et al.* (2016). Tensorflow: Large-scale machine learning on heterogeneous distributed systems. arXiv preprint arXiv:1603.04467.
- Adler, D. H. *et al.* (2018) ‘Characterizing the human hippocampus in aging and Alzheimer’s disease using a computational atlas derived from *ex-vivo* MRI and histology’, *Proceedings of the National Academy of Sciences*, pp. 4252–4257. doi: 10.1073/pnas.1801093115.
- Adnan, A. *et al.* (2016) ‘Distinct hippocampal functional networks revealed by tractography-based parcellation’, *Brain structure & function*, 221(6), pp. 2999–3012.
- Aggarwal, M. *et al.* (2015) ‘Probing region-specific microstructure of human cortical areas using high angular and spatial resolution diffusion MRI’, *NeuroImage*, pp. 198–207. doi: 10.1016/j.neuroimage.2014.10.053.
- Aggleton, J. P. (2012) ‘Multiple anatomical systems embedded within the primate medial temporal lobe: implications for hippocampal function’, *Neuroscience and biobehavioral reviews*, 36(7), pp. 1579–1596.
- Aimone, J. B., Deng, W. and Gage, F. H. (2011) ‘Resolving New Memories: A Critical Look at the Dentate Gyrus, Adult Neurogenesis, and Pattern Separation’, *Neuron*, pp. 589–596. doi: 10.1016/j.neuron.2011.05.010.
- Amunts, K. *et al.* (1999) ‘Broca’s region revisited: Cytoarchitecture and intersubject variability’, *The Journal of Comparative Neurology*, pp. 319–341. doi: 10.1002/(sici)1096-9861(19990920)412:2<319::aid-cne10>3.0.co;2-7.
- Amunts, K. *et al.* (2013) ‘BigBrain: an ultrahigh-resolution 3D human brain model’, *Science*, 340(6139), pp. 1472–1475.
- Augustinack, J. C. *et al.* (2010) ‘Direct visualization of the perforant pathway in the human brain with *ex-vivo* diffusion tensor imaging’, *Frontiers in human neuroscience*, 4, p. 42.
- Avants, B. B. *et al.* (2008) ‘Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain’, *Medical image analysis*, 12(1), pp. 26–41.
- Avants, B. B. *et al.* (2010) ‘The optimal template effect in hippocampus studies of diseased populations’, *NeuroImage*, 49(3), pp. 2457–2466.
- Avants, B. B. *et al.* (2014) ‘The Insight ToolKit image registration framework’, *Frontiers in neuroinformatics*, 8, p. 44.
- Badrinarayanan, V. *et al.* (2017) ‘SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation’, *IEEE transactions on pattern analysis and machine intelligence*, 39(12), pp. 2481–2495.

- Bakiras, S. *et al.* (2015) 'A general framework for searching in distributed data repositories', Proceedings International Parallel and Distributed Processing Symposium. doi: 10.1109/ipdps.2003.1213117.
- Berkovic, S. F. *et al.* (1991) 'Hippocampal sclerosis in temporal lobe epilepsy demonstrated by magnetic resonance imaging', *Annals of neurology*, 29(2), pp. 175–182.
- Berron, D. *et al.* (2017) 'A protocol for manual segmentation of medial temporal lobe subregions in 7 Tesla MRI', *NeuroImage. Clinical*, 15, pp. 466–482.
- Bianchi, M., Hagan, J. J. and Heidebreder, C. A. (2005) 'Neuronal plasticity, stress and depression: involvement of the cytoskeletal microtubular system?', *Current drug targets. CNS and neurological disorders*, 4(5), pp. 597–611.
- Blümcke, I. *et al.* (2007) 'A new clinico-pathological classification system for mesial temporal sclerosis', *Acta Neuropathologica*, pp. 235–244. doi: 10.1007/s00401-006-0187-0.
- Blümcke, I. *et al.* (2013) 'International consensus classification of hippocampal sclerosis in temporal lobe epilepsy: a Task Force report from the ILAE Commission on Diagnostic Methods', *Epilepsia*, 54(7), pp. 1315–1329.
- Blumcke, I. *et al.* (2017) 'Histopathological Findings in Brain Tissue Obtained during Epilepsy Surgery', *The New England journal of medicine*, 377(17), pp. 1648–1656.
- Boldrini, M. *et al.* (2013) 'Hippocampal granule neuron number and dentate gyrus volume in antidepressant-treated and untreated major depression', *Neuropsychopharmacology: official publication of the American College of Neuropsychopharmacology*, 38(6), pp. 1068–1077.
- Cai, S. *et al.* (2019) 'Is hippocampus getting bumpier with age: a quantitative analysis of fine-scale dentational feature under the hippocampus on 552 healthy subjects', *Medical Imaging 2019: Image Processing*. doi: 10.1117/12.2512701.
- Chang, C. *et al.* (2018) 'The bumps under the hippocampus', *Human brain mapping*, 39(1), pp. 472–490.
- Chappell, M. and Humphreys, M. S. (1994) 'An auto-associative neural network for sparse representations: Analysis and application to models of recognition and cued recall', *Psychological Review*, pp. 103–128. doi: 10.1037/0033-295x.101.1.103.
- Chen, L.-C. *et al.* (2018) 'DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs', *IEEE transactions on pattern analysis and machine intelligence*, 40(4), pp. 834–848.
- Cichy, R. M. *et al.* (2016) 'Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence', *Scientific reports*, 6, p. 27755.
- Clelland, C. D. *et al.* (2009) 'A functional role for adult hippocampal neurogenesis in spatial pattern separation', *Science*, 325(5937), pp. 210–213.

- Coan, A. C. *et al.* (2014) '3T MRI quantification of hippocampal volume and signal in mesial temporal lobe epilepsy improves detection of hippocampal sclerosis', *AJNR. American journal of neuroradiology*, 35(1), pp. 77–83.
- Colombo, M. *et al.* (1998) 'Functional differentiation along the anterior-posterior axis of the hippocampus in monkeys', *Journal of neurophysiology*, 80(2), pp. 1002–1005.
- Cross, J. H. *et al.* (1993) 'Early detection of abnormalities in partial epilepsy using magnetic resonance', *Archives of Disease in Childhood*, pp. 104–109. doi: 10.1136/adc.69.1.104.
- Dale, A. M., Fischl, B. and Sereno, M. I. (1999) 'Cortical surface-based analysis. I. Segmentation and surface reconstruction', *NeuroImage*, 9(2), pp. 179–194.
- de Flores, R. *et al.* (2019) 'Characterization of hippocampal subfields using *ex-vivo* MRI and histology data: Lessons for *in-vivo* segmentation', *Hippocampus*. doi: 10.1002/hipo.23172.
- de Wael, R. V., Benkarim, O., Paquola, C., Lariviere, S., Royer, J., Tavakol, S., ... & Masic, B. (2020). BrainSpace: a toolbox for the analysis of macroscale gradients in neuroimaging and connectomics datasets. *Communications biology*, 3(1), 1-10.
- DeKraker, J. *et al.* (2018) 'Unfolding the hippocampus: An intrinsic coordinate system for subfield segmentations and quantitative mapping', *NeuroImage*, 167, pp. 408–418.
- DeKraker, J. *et al.* (2020) 'Hippocampal subfields revealed through unfolding and unsupervised clustering of laminar and morphological features in 3D BigBrain', *NeuroImage*, 206, p. 116328.
- Diamond, D. M. *et al.* (1992) 'Inverted-U relationship between the level of peripheral corticosterone and the magnitude of hippocampal primed burst potentiation', *Hippocampus*, 2(4), pp. 421–430.
- Dice, L. R. (1945) 'Measures of the Amount of Ecologic Association Between Species', *Ecology*, pp. 297–302. doi: 10.2307/1932409.
- Ding, S.-L. (2013) 'Comparative anatomy of the prosubiculum, subiculum, presubiculum, postsubiculum, and parasubiculum in human, monkey, and rodent', *The Journal of comparative neurology*, 521(18), pp. 4145–4162.
- Ding, S.-L. and Van Hoesen, G. W. (2015) 'Organization and Detailed Parcellation of Human Hippocampal Head and Body Regions Based on a Combined Analysis of Cyto- and Chemoarchitecture', *The Journal of comparative neurology*, 523(15), pp. 2233–2253.
- Dinkelacker, V. *et al.* (2015) 'Hippocampal-thalamic wiring in medial temporal lobe epilepsy: Enhanced connectivity per hippocampal voxel', *Epilepsia*, 56(8), pp. 1217–1226.
- Dusek, J. A. and Eichenbaum, H. (1997) 'The hippocampus and memory for orderly stimulus relations', *Proceedings of the National Academy of Sciences of the United States of America*, 94(13), pp. 7109–7114.

- Duvernoy, H. M., Cattin, F. and Risold, P.-Y. (2013) 'The Human Hippocampus'. doi: 10.1007/978-3-642-33603-4.
- Eichenbaum, H. (2013) 'What H.M. Taught Us', *Journal of Cognitive Neuroscience*, pp. 14–21. doi: 10.1162/jocn_a_00285.
- Ekstrom, A. D. *et al.* (2009) 'Advances in high-resolution imaging and computational unfolding of the human hippocampus', *NeuroImage*, 47(1), pp. 42–49.
- Elam, J. S. and Van Essen, D. (2015) 'Human Connectome Project', *Encyclopedia of Computational Neuroscience*, pp. 1408–1411. doi: 10.1007/978-1-4614-6675-8_592.
- Engel, S. (1997) 'Retinotopic organization in human visual cortex and the spatial precision of functional MRI', *Cerebral Cortex*, pp. 181–192. doi: 10.1093/cercor/7.2.181.
- Ferko, K. (2017) 'Segmentation of Hippocampal Subfields Using 3-Tesla MRI', B. Sc Honours Specialization in Neuroscience, Co-supervised by Dr. Ali Khan and Dr Stefan Köhler, University of Western Ontario.
- Filipek, P. A. *et al.* (1994) 'The Young Adult Human Brain: An MRI-based Morphometric Analysis', *Cerebral Cortex*, pp. 344–360. doi: 10.1093/cercor/4.4.344.
- Firth, J. *et al.* (2018) 'Effect of aerobic exercise on hippocampal volume in humans: A systematic review and meta-analysis', *NeuroImage*, 166, pp. 230–238.
- Fischl, B. and Dale, A. M. (2000) 'Measuring the thickness of the human cerebral cortex from magnetic resonance images', *Proceedings of the National Academy of Sciences of the United States of America*, 97(20), pp. 11050–11055.
- Fischl, B. *et al.* (2002) 'Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain', *Neuron*, 33(3), pp. 341–355.
- Fischl, B., Sereno, M. I. and Dale, A. M. (1999) 'Cortical surface-based analysis. II: Inflation, flattening, and a surface-based coordinate system', *NeuroImage*, 9(2), pp. 195–207.
- Förster, A. *et al.* (2012) 'Diffusion-weighted imaging for the differential diagnosis of disorders affecting the hippocampus', *Cerebrovascular diseases*, 33(2), pp. 104–115.
- Frazier, J. A. *et al.* (2005) 'Structural brain magnetic resonance imaging of limbic and thalamic volumes in pediatric bipolar disorder', *The American journal of psychiatry*, 162(7), pp. 1256–1265.
- Ganepola, T. *et al.* (2018) 'Using diffusion MRI to discriminate areas of cortical grey matter', *NeuroImage*, 182, pp. 456–468.
- Garvert, M. M., Dolan, R. J. and Behrens, T. E. J. (2017) 'A map of abstract relational knowledge in the human hippocampal–entorhinal cortex', *eLife*. doi: 10.7554/elife.17086.

- Gibson, E. *et al.* (2018) ‘NiftyNet: a deep-learning platform for medical imaging’, *Computer methods and programs in biomedicine*, 158, pp. 113–122.
- Glasser, M. F. and Van Essen, D. C. (2011) ‘Mapping Human Cortical Areas In Vivo Based on Myelin Content as Revealed by T1- and T2-Weighted MRI’, *Journal of Neuroscience*, pp. 11597–11616. doi: 10.1523/jneurosci.2180-11.2011.
- Glasser, M. F. *et al.* (2013) ‘The minimal preprocessing pipelines for the Human Connectome Project’, *NeuroImage*, 80, pp. 105–124.
- Glasser, M. F. *et al.* (2014) ‘Trends and properties of human cerebral cortex: Correlations with cortical myelin content’, *NeuroImage*, pp. 165–175. doi: 10.1016/j.neuroimage.2013.03.060.
- Glasser, M. F. *et al.* (2016) ‘A multi-modal parcellation of human cerebral cortex’, *Nature*, pp. 171–178. doi: 10.1038/nature18933.
- Goodfellow, I., Bengio, Y. and Courville, A. (2016) *Deep Learning*. MIT Press.
- Goubran, M. *et al.* (2013) ‘Image registration of *ex-vivo* MRI to sparsely sectioned histology of hippocampal and neocortical temporal lobe specimens’, *NeuroImage*, pp. 770–781. doi: 10.1016/j.neuroimage.2013.07.053.
- Goubran, M. *et al.* (2015) ‘Registration of *in-vivo* to *ex-vivo* MRI of surgically resected specimens: A pipeline for histology to *in-vivo* registration’, *Journal of Neuroscience Methods*, pp. 53–65. doi: 10.1016/j.jneumeth.2014.12.005.
- Gould, E. (1999) ‘Serotonin and Hippocampal Neurogenesis’, *Neuropsychopharmacology*, pp. S46–S51. doi: 10.1038/sj.npp.1395369.
- Graving, J. M. *et al.* (2019) ‘DeepPoseKit, a software toolkit for fast and robust animal pose estimation using deep learning’, *eLife*, 8. doi: 10.7554/eLife.47994.
- Gross, D. W. *et al.* (2020) ‘Curved multiplanar reformatting provides improved visualization of hippocampal anatomy’, *Hippocampus*, 30(2), pp. 156–161.
- Hashemi, R. H., Bradley, W. G., & Lisanti, C. J. (2012). *MRI: the basics: The Basics*. Lippincott Williams & Wilkins.
- He, K. *et al.* (2016) ‘Deep Residual Learning for Image Recognition’, 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). doi: 10.1109/cvpr.2016.90.
- Henry, T. R. *et al.* (2011) ‘Hippocampal Sclerosis in Temporal Lobe Epilepsy: Findings at 7 T’, *Radiology*, pp. 199–209. doi: 10.1148/radiol.11101651.
- Heras, F. J. H. *et al.* (2019) ‘Deep attention networks reveal the rules of collective motion in zebrafish’, *PLoS computational biology*, 15(9), p. e1007354.
- Hindy, N. C., Ng, F. Y. and Turk-Browne, N. B. (2016) ‘Linking pattern completion in the hippocampus to predictive coding in visual cortex’, *Nature neuroscience*, 19(5), pp. 665–667.
- Hunsaker, M. R. and Kesner, R. P. (2013) ‘The operation of pattern separation and pattern completion processes associated with different attributes or domains of

- memory', *Neuroscience & Biobehavioral Reviews*, pp. 36–58. doi: 10.1016/j.neubiorev.2012.09.014.
- Hussain, U. *et al.* (in preparation) 'Diffusion MRI of the Unfolded Hippocampus'. doi: 10.1101/2020.04.07.029850.
- Iglesias, J. E. *et al.* (2015) 'A computational atlas of the hippocampal formation using ex vivo, ultra-high resolution MRI: Application to adaptive segmentation of *in-vivo* MRI', *NeuroImage*, 115, pp. 117–137.
- Jackson, G. D. *et al.* (1990) 'Hippocampal sclerosis can be reliably detected by magnetic resonance imaging', *Neurology*, 40(12), pp. 1869–1875.
- Jenkinson, M. and Smith, S. (2001) 'A global optimisation method for robust affine registration of brain images', *Medical image analysis*, 5(2), pp. 143–156.
- Jenkinson, M. *et al.* (2002) 'Improved optimization for the robust and accurate linear registration and motion correction of brain images', *NeuroImage*, 17(2), pp. 825–841.
- Jones, S. E., Buchbinder, B. R. and Aharon, I. (2000) 'Three-dimensional mapping of cortical thickness using Laplace's Equation', *Human Brain Mapping*, pp. 12–32. doi: 10.1002/1097-0193(200009)11:1<12::aid-hbm20>3.0.co;2-k
- Kerchner, G. A. *et al.* (2010) 'Hippocampal CA1 apical neuropil atrophy in mild Alzheimer disease visualized with 7-T MRI', *Neurology*, 75(15), pp. 1381–1387.
- Kerchner, G. A. *et al.* (2012) 'Hippocampal CA1 apical neuropil atrophy and memory performance in Alzheimer's disease', *NeuroImage*, 63(1), pp. 194–202.
- Kim, J. S. *et al.* (2005) 'Automated 3-D extraction and evaluation of the inner and outer cortical surfaces using a Laplacian map and partial volume effect classification', *NeuroImage*, 27(1), pp. 210–221.
- Knierim, J. J. and Zhang, K. (2012) 'Attractor dynamics of spatially correlated neural activity in the limbic system', *Annual review of neuroscience*, 35, pp. 267–285.
- Kozareva, D. A., Cryan, J. F. and Nolan, Y. M. (2019) 'Born this way: Hippocampal neurogenesis across the lifespan', *Aging cell*, 18(5), p. e13007.
- Krizhevsky, A., Sutskever, I. and Hinton, G. E. (2017) 'ImageNet classification with deep convolutional neural networks', *Communications of the ACM*, pp. 84–90. doi: 10.1145/3065386.
- Kyle, C. T. *et al.* (2015) 'Successful retrieval of competing spatial environments in humans involves hippocampal pattern separation mechanisms', *eLife*, 4. doi: 10.7554/eLife.10499.
- Ladd, M. E. *et al.* (2018) 'Pros and cons of ultra-high-field MRI/MRS for human application', *Progress in nuclear magnetic resonance spectroscopy*, 109, pp. 1–50.
- Landman, B. A. *et al.* (2013) 'System for integrated neuroimaging analysis and processing of structure', *Neuroinformatics*, 11(1), pp. 91–103.
- Larsen, R., Nielsen, M. and Sparring, J. (2006) *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2006: 9th International Conference*,

Copenhagen, Denmark, October 1-6, 2006, Proceedings. Springer Science & Business Media.

- Lauer, S. M. *et al.* (2018) ‘Visualization of Cortical Modules in Flattened Mammalian Cortices’, *Journal of Visualized Experiments*. doi: 10.3791/56992.
- Leal, S. L. and Yassa, M. A. (2018) ‘Integrating new findings and examining clinical applications of pattern separation’, *Nature neuroscience*, 21(2), pp. 163–173.
- Lee, A. C. H., Yeung, L.-K. and Barense, M. D. (2012) ‘The hippocampus and visual perception’, *Frontiers in Human Neuroscience*. doi: 10.3389/fnhum.2012.00091.
- Lehéricy, S. *et al.* (1995) ‘Developmental abnormalities of the medial temporal lobe in patients with temporal lobe epilepsy’, *AJNR. American journal of neuroradiology*, 16(4), pp. 617–626.
- Leuze, C. W. U. *et al.* (2014) ‘Layer-specific intracortical connectivity revealed with diffusion MRI’, *Cerebral cortex*, 24(2), pp. 328–339.
- Li, W. *et al.* (2017) ‘On the Compactness, Efficiency, and Representation of 3D Convolutional Networks: Brain Parcellation as a Pretext Task’, *Lecture Notes in Computer Science*, pp. 348–360. doi: 10.1007/978-3-319-59050-9_28.
- Li, X.-W. *et al.* (2015) ‘Local manifold learning for multiatlas segmentation: application to hippocampal segmentation in healthy population and Alzheimer’s disease’, *CNS neuroscience & therapeutics*, 21(10), pp. 826–836.
- Liu, K. Y. *et al.* (2016) ‘Tests of pattern separation and pattern completion in humans-A systematic review’, *Hippocampus*, 26(6), pp. 705–717.
- Lubenov, E. V., & Siapas, A. G. (2009). Hippocampal theta oscillations are travelling waves. *Nature*, 459(7246), 534-539.
- MacDonald, D. *et al.* (2000) ‘Automated 3-D extraction of inner and outer surfaces of cerebral cortex from MRI’, *NeuroImage*, 12(3), pp. 340–356.
- Madison, D. V. and Nicoll, R. A. (1982) ‘Noradrenaline blocks accommodation of pyramidal cell discharge in the hippocampus’, *Nature*, 299(5884), pp. 636–638.
- Malenka, R. C. and Bear, M. F. (2004) ‘LTP and LTD: an embarrassment of riches’, *Neuron*, 44(1), pp. 5–21.
- Malykhin, N. V. and Coupland, N. J. (2015) ‘Hippocampal neuroplasticity in major depressive disorder’, *Neuroscience*, pp. 200–213. doi: 10.1016/j.neuroscience.2015.04.047.
- Marr, D. (1971) ‘Simple memory: a theory for archicortex’, *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 262(841), pp. 23–81.
- Maruszak, A. and Thuret, S. (2014) ‘Why looking at the whole hippocampus is not enough-a critical role for anteroposterior axis, subfield and activation analyses to enhance predictive value of hippocampal changes for Alzheimer’s disease diagnosis’, *Frontiers in cellular neuroscience*, 8, p. 95.

- Mathis, A. *et al.* (2018) ‘DeepLabCut: markerless pose estimation of user-defined body parts with deep learning’, *Nature Neuroscience*, pp. 1281–1289. doi: 10.1038/s41593-018-0209-y.
- McClelland, J. L., McNaughton, B. L. and O’Reilly, R. C. (1995) ‘Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory’, *Psychological Review*, pp. 419–457. doi: 10.1037/0033-295x.102.3.419.
- Menschik, E. D. and Finkel, L. H. (1998) ‘Neuromodulatory control of hippocampal function: towards a model of Alzheimer’s disease’, *Artificial Intelligence in Medicine*, pp. 99–121. doi: 10.1016/s0933-3657(98)00006-2.
- Mitchell, D. G. and Mark Cohen (ph (2004) *MRI Principles*. W B Saunders Company.
- Moscovitch, M. *et al.* (2016) ‘Episodic Memory and Beyond: The Hippocampus and Neocortex in Transformation’, *Annual review of psychology*, 67, pp. 105–134.
- Mueller, S. G. *et al.* (2008) ‘Selective effect of Apo e4 on CA3 and dentate in normal aging and Alzheimer’s disease using high resolution MRI at 4 T’, *NeuroImage*, 42(1), pp. 42–48.
- Murray, E. A., Bussey, T. J. and Saksida, L. M. (2007) ‘Visual perception and memory: a new view of medial temporal lobe function in primates and rodents’, *Annual review of neuroscience*, 30, pp. 99–122.
- Na, K.-S. *et al.* (2014) ‘Association between glucocorticoid receptor methylation and hippocampal subfields in major depressive disorder’, *PloS one*, 9(1), p. e85425.
- Nath, T. *et al.* (2018) ‘Using DeepLabCut for 3D markerless pose estimation across species and behaviors’. doi: 10.1101/476531.
- Nieuwenhuys, R., Voogd, J. and van Huijzen, C. (1981) ‘The Human Central Nervous System’. doi: 10.1007/978-3-662-02333-4.
- Nordanskog, P. *et al.* (2010) ‘Increase in hippocampal volume after electroconvulsive therapy in patients with depression: a volumetric magnetic resonance imaging study’, *The journal of ECT*, 26(1), pp. 62–67.
- Norman, K.A. and O’Reilly, R.C., 2003. Modeling hippocampal and neocortical contributions to recognition memory: a complementary-learning-systems approach. *Psychological review*, 110(4), p.611.
- O’Keefe, J., Nadel, L. (1978) *The Hippocampus as a Cognitive Map*. Oxford University Press, USA.
- Olah, C. *et al.* (2018) ‘The Building Blocks of Interpretability’, *Distill*. doi: 10.23915/distill.00010.
- Olsen, R. K. *et al.* (2019) ‘Progress update from the hippocampal subfields group’, *Alzheimer’s & dementia: the journal of the Alzheimer’s Association*, 11, pp. 439–449.

- Oppenheim, C. *et al.* (1999) ‘Bilateral mesial temporal sclerosis: MRI with high-resolution fast spin-echo and fluid-attenuated inversion-recovery sequences’, *Neuroradiology*, pp. 471–479. doi: 10.1007/s002340050786.
- Oppenheim, C. *et al.* (1998). Loss of digitations of the hippocampal head on high-resolution fast spin-echo MR: a sign of mesial temporal sclerosis. *American journal of neuroradiology*, 19(3), pp.457-463.
- Paquola, C. *et al.* (in submission) ‘Convergence of cortical types and functional motifs in the mesiotemporal lobe’. doi: 10.1101/2020.06.12.148643.
- Pereira, J. B. *et al.* (2014) ‘Regional vulnerability of hippocampal subfields to aging measured by structural and diffusion MRI’, *Hippocampus*, 24(4), pp. 403–414.
- Plachti, A. *et al.* (2019) ‘Multimodal Parcellations and Extensive Behavioral Profiling Tackling the Hippocampus Gradient’, *Cerebral cortex*, 29(11), pp. 4595–4612.
- Poppenk, J. *et al.* (2013) ‘Long-axis specialization of the human hippocampus’, *Trends in cognitive sciences*, 17(5), pp. 230–240.
- Possin, K. L. (2010) ‘Visual spatial cognition in neurodegenerative disease’, *Neurocase*, pp. 466–487. doi: 10.1080/13554791003730600.
- Pravosudov, V. V. and Roth, T. C., II (2013) ‘Cognitive Ecology of Food Hoarding: The Evolution of Spatial Memory and the Hippocampus’, *Annual Review of Ecology, Evolution, and Systematics*, pp. 173–193. doi: 10.1146/annurev-ecolsys-110512-135904.
- Puelles, L. *et al.* (2019) ‘Concentric ring topology of mammalian cortical sectors and relevance for patterning studies’, *The Journal of comparative neurology*, 527(10), pp. 1731–1752.
- Qiu, A. and Miller, M. I. (2008) ‘Multi-structure network shape analysis via normal surface momentum maps’, *NeuroImage*, 42(4), pp. 1430–1438.
- Research, C. M. and Case Medical Research (2019) ‘Fully Automated Pipeline for the Detection and Segmentation of Non-Small Cell Lung Cancer (NSCLC) on CT Images’, *Case Medical Research*. doi: 10.31525/ct1-nct04164186.
- Rolls, E. T. (2016) ‘Pattern separation, completion, and categorisation in the hippocampus and neocortex’, *Neurobiology of learning and memory*, 129, pp. 4–28.
- Rolls, E. T. and Kesner, R. P. (2016) ‘Pattern separation and pattern completion in the hippocampal system. Introduction to the Special Issue’, *Neurobiology of learning and memory*, 129, pp. 1–3.
- Ronneberger, O., Fischer, P. and Brox, T. (2015) ‘U-Net: Convolutional Networks for Biomedical Image Segmentation’, *Lecture Notes in Computer Science*, pp. 234–241. doi: 10.1007/978-3-319-24574-4_28.
- Rumelhart, D. E., Hinton, G. E. and Williams, R. J. (1986) ‘Learning representations by back-propagating errors’, *Nature*, pp. 533–536. doi: 10.1038/323533a0.

- Sahay, A. and Hen, R. (2007) 'Adult hippocampal neurogenesis in depression', *Nature Neuroscience*, pp. 1110–1115. doi: 10.1038/nn1969.
- Save, E. and Poucet, B. (2000) 'Hippocampal-parietal cortical interactions in spatial cognition', *Hippocampus*, 10(4), pp. 491–499.
- Schacter, D. L., Benoit, R. G. and Szpunar, K. K. (2017) 'Episodic Future Thinking: Mechanisms and Functions', *Current opinion in behavioral sciences*, 17, pp. 41–50.
- Schapiro, A. C., Turk-Browne, N. B., Botvinick, M. M., & Norman, K. A. (2017). Complementary learning systems within the hippocampus: a neural network modelling approach to reconciling episodic memory with statistical learning. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372(1711), 20160049.
- Schleicher, A. *et al.* (1999) 'Observer-independent method for microstructural parcellation of cerebral cortex: A quantitative approach to cytoarchitectonics', *NeuroImage*, 9(1), pp. 165–177.
- Schmidt, M. F. *et al.* (2018) 'A comparison of manual tracing and FreeSurfer for estimating hippocampal volume over the adult lifespan', *Human brain mapping*, 39(6), pp. 2500–2513.
- Sethian, J. A. (1996) *Level Set Methods: Evolving Interfaces in Geometry, Fluid Mechanics, Computer Vision, and Materials Science*.
- Shi, Y., Cheng, K. and Liu, Z. (2019) 'Hippocampal subfields segmentation in brain MR images using generative adversarial networks', *Biomedical engineering online*, 18(1), p. 5.
- Small, S. A. *et al.* (2011) 'A pathophysiological framework of hippocampal dysfunction in ageing and disease', *Nature Reviews Neuroscience*, pp. 585–601. doi: 10.1038/nrn3085.
- Sørensen, T. (1948) *A Method of Establishing Groups of Equal Amplitude in Plant Sociology Based on Similarity of Species Content and Its Application to Analyses of the Vegetation on Danish Commons*.
- Sowell, E. R. (2004) 'Longitudinal Mapping of Cortical Thickness and Brain Growth in Normal Children', *Journal of Neuroscience*, pp. 8223–8231. doi: 10.1523/jneurosci.1798-04.2004.
- Sperber, G. H. (1995) 'Developmental biology, 4th Edition. By Scott F. Gilbert, Sinauer Associates, Inc., Sunderland, Massachusetts, and W.H. Freeman, U.K., 1994, pp. xviii, 894, \$57.95/£29.95', *American Journal of Medical Genetics*, pp. 642–642. doi: 10.1002/ajmg.1320570429.
- Squire, L. R. (2009) 'The Legacy of Patient H.M. for Neuroscience', *Neuron*, pp. 6–9. doi: 10.1016/j.neuron.2008.12.023.
- Steve, T. A. *et al.* (2020) 'Hippocampal subfield measurement and ILAE hippocampal sclerosis subtype classification with *in-vivo* 4.7 tesla MRI', *Epilepsy research*, 161, p. 106279.

- Strange, B. A. *et al.* (2014) 'Functional organization of the hippocampal longitudinal axis', *Nature reviews. Neuroscience*, 15(10), pp. 655–669.
- Sudre, C. H., Li, W., Vercauteren, T., Ourselin, S., & Cardoso, M. J. (2017). Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In *Deep learning in medical image analysis and multimodal learning for clinical decision support* (pp. 240-248). Springer, Cham.
- Treit, S. *et al.* (2018) 'High resolution *in-vivo* diffusion imaging of the human hippocampus', *NeuroImage*, 182, pp. 479–487.
- Tsai, M.-H. *et al.* (2016) 'Hippocampal malrotation is an anatomic variant and has no clinical significance in MRI-negative temporal lobe epilepsy', *Epilepsia*, 57(10), pp. 1719–1728.
- Tsao, J. and Kozerke, S. (2012) 'MRI temporal acceleration techniques', *Journal of Magnetic Resonance Imaging*, pp. 543–560. doi: 10.1002/jmri.23640.
- Tulving, E. and Markowitsch, H. J. (1998) 'Episodic and declarative memory: Role of the hippocampus', *Hippocampus*, pp. 198–204. doi: 10.1002/(sici)1098-1063(1998)8:3<198::aid-hipo2>3.0.co;2-g.
- Tustison, N. J. *et al.* (2014) 'Large-scale evaluation of ANTs and FreeSurfer cortical thickness measurements', *NeuroImage*, 99, pp. 166–179.
- Tyszka, J. M., Michael Tyszka, J. and Pauli, W. M. (2016) 'In vivo delineation of subdivisions of the human amygdaloid complex in a high-resolution group template', *Human Brain Mapping*, pp. 3979–3998. doi: 10.1002/hbm.23289.
- Van Leemput, K. *et al.* (2009) 'Automated segmentation of hippocampal subfields from ultra-high resolution *in-vivo* MRI', *Hippocampus*, 19(6), pp. 549–557.
- Vos de Wael, R. *et al.* (2018) 'Anatomical and microstructural determinants of hippocampal subfield functional connectome embedding', *Proceedings of the National Academy of Sciences of the United States of America*, 115(40), pp. 10154–10159.
- Waehnert, M. D. *et al.* (2014) 'Anatomically motivated modeling of cortical laminae', *NeuroImage*, pp. 210–220. doi: 10.1016/j.neuroimage.2013.03.078.
- Wagstyl, K. *et al.* (2018) 'Mapping Cortical Laminar Structure in the 3D BigBrain', *Cerebral cortex*, 28(7), pp. 2551–2562.
- Winterburn, J. L. *et al.* (2013) 'A novel *in-vivo* atlas of human hippocampal subfields using high-resolution 3 T magnetic resonance imaging', *NeuroImage*, 74, pp. 254–265.
- Wisse, L. E. M. *et al.* (2012) 'Subfields of the hippocampal formation at 7T MRI: In vivo volumetric assessment', *NeuroImage*, pp. 1043–1049. doi: 10.1016/j.neuroimage.2012.03.023.
- Wisse, L. E. M. *et al.* (2016) 'Automated Hippocampal Subfield Segmentation at 7T MRI', *AJNR. American journal of neuroradiology*, 37(6), pp. 1050–1057.

- Wisse, L. E. M. *et al.* (2017) ‘A harmonized segmentation protocol for hippocampal and parahippocampal subregions: Why do we need one and what are the key goals?’, *Hippocampus*, 27(1), pp. 3–11.
- Wisse, L. E. M., Biessels, G. J. and Geerlings, M. I. (2014) ‘A Critical Appraisal of the Hippocampal Subfield Segmentation Package in FreeSurfer’, *Frontiers in aging neuroscience*, 6, p. 261.
- Wolz, R. *et al.* (2010) ‘LEAP: learning embeddings for atlas propagation’, *NeuroImage*, 49(2), pp. 1316–1325.
- Xie, L. *et al.* (2017) ‘Multi-template analysis of human perirhinal cortex in brain MRI: Explicitly accounting for anatomical variability’, *NeuroImage*, 144(Pt A), pp. 183–202.
- Yang, Z. *et al.* (2020) ‘CAST: A multi-scale convolutional neural network based automated hippocampal subfield segmentation toolbox’, *NeuroImage*, p. 116947. doi: 10.1016/j.neuroimage.2020.116947.
- Yassa, M. A., Muftuler, L. T. and Stark, C. E. L. (2010) ‘Ultra-high-resolution microstructural diffusion tensor imaging reveals perforant path degradation in aged humans *in vivo*’, *Proceedings of the National Academy of Sciences of the United States of America*, 107(28), pp. 12687–12691.
- Ye, J. C. (2019) ‘Compressed sensing MRI: a review from signal processing perspective’, *BMC Biomedical Engineering*. doi: 10.1186/s42490-019-0006-z.
- Yushkevich, P. A. *et al.* (2006) ‘User-guided 3D active contour segmentation of anatomical structures: significantly improved efficiency and reliability’, *NeuroImage*, 31(3), pp. 1116–1128.
- Yushkevich, P. A. *et al.* (2015) ‘Quantitative comparison of 21 protocols for labeling hippocampal subfields and parahippocampal subregions in *in-vivo* MRI: Towards a harmonized segmentation protocol’, *NeuroImage*, pp. 526–541. doi: 10.1016/j.neuroimage.2015.01.004.
- Yushkevich, P. A., Pluta, J. B., *et al.* (2015) ‘Automated volumetry and regional thickness analysis of hippocampal subfields and medial temporal cortical structures in mild cognitive impairment’, *Human brain mapping*, 36(1), pp. 258–287.
- Zaeemzadeh, A., Rahnavard, N. and Shah, M. (2020) ‘Norm-Preservation: Why Residual Networks Can Become Extremely Deep?’, *IEEE transactions on pattern analysis and machine intelligence*. doi: 10.1109/TPAMI.2020.2990339.
- Zeidman, P. and Maguire, E. A. (2016) ‘Anterior hippocampus: the anatomy of perception, imagination and episodic memory’, *Nature reviews. Neuroscience*, 17(3), pp. 173–182.
- Zeineh, M. M. *et al.* (2001) ‘Unfolding the human hippocampus with high resolution structural and functional MRI’, *The Anatomical record*, 265(2), pp. 111–120.
- Zhu, H. *et al.* (2019) ‘Dilated Dense U-Net for Infant Hippocampus Subfield Segmentation’, *Frontiers in neuroinformatics*, 13, p. 30.

Zijdenbos, A. P., Forghani, R. and Evans, A. C. (2002) 'Automatic "pipeline" analysis of 3-D MRI data for clinical trials: application to multiple sclerosis', *IEEE transactions on medical imaging*, 21(10), pp. 1280–1291.

Appendices

Appendix 1 (Chapter 2): Manual and user-guided steps for labelling of SRLM, hippocampal grey matter, and extrahippocampal structures for subsequent unfolding

Criteria for automatic unfolding under current code setup:

- Hippocampus should be split into 3 labels in the hippocampus: SRLM, hippocampal grey matter, and cysts
- Four extra-hippocampal structure labels are needed for subsequent unfolding: border with the medial-temporal lobe cortex (MTLc), medial side of vertical component of the uncus (mVUnc) (for determining thickness in this area; also serves to mark dentate gyrus in this area), hippocampal-amygdalar transition area (HATA), and indusium griseum.
- SRLM should have no 'holes' in it (e.g. when viewed in 3D with no model smoothing).
- SRLM should be visible on the finished 3D model all along its medial edge (above subiculum) and nowhere else (i.e. no 'holes' in grey matter label).
- Hippocampal grey matter label should not contain or be touching any blank label except on its outer surface (i.e. no blank voxels inside hippocampus)
- Cyst label should only include CSF inside the hippocampus. If there are cysts that pass to the outside of the hippocampus, they should be ignored.
- Extra-hippocampal structures should be labelled where they border the hippocampus (their entire structure can be labeled, but it's not necessary performing subsequent unfolding).
- Consideration of folding in a manual segmentation is recommended. If hippocampal grey matter from one fold contacts a different fold (e.g. in the digitations or along the curve of the uncus), then the automatic unfolding of the hippocampus will be incorrect (think of the Laplacian filter as 'leaking' through holes in the SRLM or across overlapping folds).
- We HIGHLY RECOMMEND viewing the finished segmentation models we have provided in addition to the instructions here. Even better is to keep the model that most closely resembles a segmentation being performed open in another window, as a reference.

1. Tracing of the SRLM

The SRLM separates the dentate gyrus on the inside from the outer subfields, and therefore the SRLM should always be some distance inward from the outer

hippocampal border, surrounded by hippocampal grey matter. The exception to this rule is along the medial side of the hippocampus, where the SRLM extends over subiculum to the ventricle. The SRLM may sometimes be obscured by hippocampal cysts, in which case it should be estimated and drawn in or the cyst label should completely 'plug' the hole in SRLM label. It's very easy to confuse SRLM with alveus, particularly in the axial and sagittal planes. Thus, only dark areas which clearly constitute SRLM should be traced on the first pass, and missing areas can be filled in later. Where uncertain, use of the other planes of view is recommended.

We have found tracing with a 2x2 paintbrush to offer good feasibility.

The label for the SRLM should be saved separately from the labels which will later be generated from it.




Fig 1.1. Models of hippocampal SRLM in a highly digitated (left) and a less digitated hippocampus (right). Green is labelled SRLM and yellow is hippocampal cysts.

Axial view. In this view, at the approximate middle of the hippocampus, the hippocampus should look like a footprint, the anterior digitations being the toes. The SRLM is just inward of the outer border, and should be traced only when it is clear, and only when it passes through the plane (as opposed to parts where it runs parallel to this plane, which will appear as larger, poorly defined dark blobs). Starting in the middle of the hippocampus and working downward and then upward is recommended.




Fig 1.2. Axial view of SRLM trace. Green is labelled SRLM and yellow is hippocampal cysts.

Sagittal view. The same rules as 1.2 apply here. There is no SRLM over top of most of the subiculum in more medial views. This view should be used for tracing the hippocampal tail, which should be included as far posteriorly and medially as the SRLM is obvious. Dilation of the SRLM in the next step should include all the grey matter that touches the hippocampus, even if it appears to extend somewhat far posteriorly. It may be necessary to relabel this far posterior grey matter as indusium griseum later when labelling extra-hippocampal structures.




Fig 1.3. Sagittal view of SRLM on the lateral side (top) and very medial side where only

hippocampal tail is present (bottom). Green is labelled SRLM and yellow is hippocampal cysts

Coronal view. This is where the SRLM should be most clear in many cases and should have a distinctive 'C' shape throughout the hippocampal body. In the head this shape is modified to extend into each digitation, and at the very most anterior the 'C' closes in on itself and becomes a single line. Note that at the superior side, the surface subfield CA3 often appears dark and so great care should be taken to separate SRLM from CA3. After tracing what is visible in all slices, it should become clear from the model what parts of SRLM were missed (This is all one continuous tissue! It is critical to look for the holes). Any large gaps should be filled in, and when uncertain about how a piece of SRLM connects to the rest of the SRLM structure, the other planes and the example segmentations provided should be considered. After hippocampal grey matter is traced, the SRLM will still border the medial side of the hippocampus along the entire anterior-posterior extent except at the most anterior tip, around where the VUnc (which will be traced in the next steps) is no longer visible. Looking through this opening, the concavity that the SRLM forms enclosing the dentate gyrus should be visible. It is important to track tissues across different slices or across the different planes to determine what they are and to confirm that labels line up correctly.




Fig 1.4. Coronal views of hippocampal SRLM from body (top), through the digitations in the hippocampal head, to the anterior where it flattens out into a single line (descending). Green is labelled SRLM and yellow is hippocampal cysts

Touch-ups. A good way to smooth SRLM label and also fill in any small gaps is to dilate and then contract the entire label by a few voxels, a technique known as ‘feathering’. This can be done using the ‘Snake’ tool in ITK-SNAP – use of the edge attraction operation is recommended, with the filter options such that no edges are visible. Then roughly 2-10 iterations of positive growth should be run, followed by the same number of iterations of negative growth. It is important not to grow the label so much that different folds become connected!

When confident, tracing of every other slice may be sufficient in order to save time. The feathering of the label will fill the gaps in between. Manual adjustments to the SRLM can be made later too, so if uncertain then go to the next step.

It is important to try to keep the SRLM label thin! Contracting the label as far as possible without opening up gaps will be helpful to best capture only the dark voxels in the SRLM.

2. Generation of hippocampal grey matter label

The SRLM label contains most of the spatial information we need – it enters into each digitation and runs parallel to the inferior, lateral, and much of the superior edge of the hippocampus (the medial and superior edges will still need some work).




Fig 2.1. Whole hippocampus model, with the SRLM visible in green along the medial side. Green is labelled SRLM, red is hippocampal grey matter, and yellow is hippocampal cysts.

Dilation of the SRLM. This can be done in ITK-SNAP using the Snake tool. I’ve found the best way to do this is run unconstrained dilation (i.e. no thresholding or edge attraction) to almost the outer edge of hippocampal grey matter, and then introduce some edge attraction constraints for the rest of the dilation. This avoids attraction to other features, but still allows grey matter thickness to vary slightly. The goal here is to best fit as many of the outer edges of the hippocampus as possible. Some edges, particularly

the medial, will have to be manually adjusted no matter what criteria are set, so focus on the lateral, superior, and inferior borders.

This should be saved separately from the SRLM label. Next the SRLM label should be combined with this dilated label (for example, add the two label images together using FSL's `fslmaths` tool). The resulting label should contain 1's for whole hippocampus and 2's for SRLM.




Figure 2.2. Dilation of SRLM in the body (top) and head (bottom).

Manual adjustments. Along the superior edge, the dilation may have gone past the border of the hippocampus since the subfield present there – the CA3 – is typically thin. Thus, it will have to be trimmed down, and may have to be extended slightly medially to fill in any additional grey matter that was missed. Similarly, dark blood vessels, alveus, and CSF are also often present medially to the dentate gyrus and superior to the subiculum and should be removed (see Fig 2.3).

Some patches in the center of the dentate gyrus may have been missed, and some minor adjustments to the lateral and inferior borders are sometimes necessary. Additionally, the alveus entering in between each digitation is sometimes filled in because of too much dilation. This should also be manually removed such that label from one digitation is separated from other digitations by at least 1 voxel in all orientations (including diagonally). SRLM can also be manually adjusted while performing these adjustments.

Looking at the model, the SRLM should be cleanly visible on the medial side. In the uncus, the SRLM should still be visible and sometimes a large area of SRLM label is visible on the model below VUnc. In these cases, it is critical to remove all of the dilated hippocampal label medial to the SRLM. The grey matter of the uncus should not contact the grey matter of the subiculum running below it - they should be separated by clear label or otherwise by SRLM.

Because the hippocampal tail curves medially, the SRLM will be visible on the anterior side.

The thresholding tool should be used across the entire hippocampus to relabel any cysts or CSF around the hippocampus as such. Look for round, light 'patches' that fade gradually into more typical grey matter and exclude the light parts. If this does not work well, these cysts can simply be labelled manually (we recommend a spherical, 4-voxel diameter 3D brush). Where CSF is visible in one slice, there is also likely to be at least some partial voluming with CSF in the next slices.




Fig 2.3. Examples of manual adjustments to be made to the whole hippocampus label in the body (top) and head (bottom). Green is labelled SRLM, red is hippocampal grey matter, and yellow is hippocampal cysts.

Subiculum. The subiculum is typically dark and is sometimes almost indistinguishable

from white matter. However, even if it is difficult to distinguish, a band of tissue (~1mm in thickness) should extend medially from the hippocampus towards MTLc. This tissue runs below the CSF in the ventricle and should not include partial voluming from the ventricle. The most medial border of subiculum is defined here as the most medial extension of white matter (though note that in Ding *et al.* (2016), parasubiculum sometimes extends slightly past this point, which we did not include here for simplicity). In the hippocampal head, the subiculum runs almost entirely inferior to the SRLM and so not much will need to be added (though there should still be a small 'lip' visible extending downward and medially. In the hippocampal tail, we use the following rule: the medial extension that is subiculum reaches the medial extent of white, which becomes shorter until it collides with the most medial, posterior edge of the hippocampal tail (see Fig 2.1) (though note that according to Ding *et al.*, (2016) this tissue also contains some retrosplenial area 29).




Fig 2.4. Tracing of medial extension making up the subiculum in the body (top) and head (bottom). Green is labelled SRLM, red is hippocampal grey matter, and yellow is hippocampal cysts.

Vertical component of the uncus. This must be drawn in manually (or carefully grown in using the Snake tool). It is a thin layer of grey matter on the medial side of the uncus and separated from the rest of the uncus by white matter of the alveus. Often there is CSF medial to this, so care is required not to include partial volume CSF. In anterior

slices the superior border is not clear and should be drawn approximately at the most superior part of the rest of the hippocampus. Past this is defined here as hippocampal-amygdala transition area (though note that according to Ding *et al.* (2016) this border is sometimes even more superior). The VUnc should become shorter in the most anterior until it is indistinguishable from the rest of the hippocampal head at roughly the same coronal slice as the SRLM can no longer be seen on the medial side of the 3D model.




Fig 2.5. Tracing of vertical component of the uncus (VUnc). Green is labelled SRLM, red is hippocampal grey matter, and yellow is hippocampal cysts.

3. Labeling of extra-hippocampal structures

These labels aren't part of the hippocampus, but are needed to mark the anterior, posterior, and medial edges of the hippocampus in subsequent unfolding. Since these aren't the primary structures of interest here, we typically label them only where they actually border the hippocampus. Each of these labels should contact one of the others (e.g. HATA touches MTLc label - so subsequent unfolding will be close to orthogonal).




Fig 3.1. Extra-hippocampal structures in the head (top) and tail (bottom). Green is labelled SRLM and red is hippocampal grey matter. The MTLc border is marked in white, HATA in pink, indusium griseum in purple, and the mVUnc in yellow. Labels include only parts of their structure that border the hippocampus.

Medial-temporal lobe cortex border. This should border the medial edge of the subiculum. In the hippocampal head, it wraps upward and medially along with the subiculum, and terminates on the VUnc (also where the hippocampal-amygdalar transition area ends). Note that the MTLc label should not touch the inferior-posterior side of the uncus (where it runs over the subiculum in the hippocampal head). In the hippocampal tail this border should extend back until it contacts the indusium griseum label.

Hippocampal-Amygdala Transition Area (HATA). This should be marked on the grey matter on the superior edge of the VUnc, just below and medial to the amygdala. The cutoff we used to distinguish VUnc from HATA was a straight, horizontal line at the superior most extent of the rest of the hippocampal head. Its posterior border is where no grey matter 'bridge' can be seen connecting the VUnc to the amygdala, and its anterior border is where the VUnc is no longer visible and is instead replaced with MTLc. The superior border is where this grey matter bridge abruptly widens, forming the inferior-medial edge of the amygdala. However, this superior border was not always enforced in our tracing as we were only interested in the components of HATA that directly bordered the hippocampus.




Fig 3.2. Borders of the HATA. First row shows the anterior border where no VUnc is visible. Second row shows the bridge between hippocampus and amygdala, and the third row shows the posterior border where this bridge is no longer visible (blue arrows). Blue line shows the border between hippocampal tissue and HATA at the superior extent of hippocampus.

Medial edge of vertical component of the uncus (mVUnc). This label is a dummy label, defined geometrically in order to mark the medial edge of this tissue as separate from the lateral edge when determining cortical thickness. It should completely cover the hippocampal grey matter label where the VUnc appears as a single, thin layer of tissue.

The Dentate Gyrus granule cell layer (DGgcl) within the vertical component of the

uncus, which is used as the sink in the process of proximal-distal unfolding, is also derived from this label: the posterior edge of mVUnc will mark where the dentate gyrus extends through the VUnc on the medial side. Thus, the posterior border of mvUnc should be the most posterior third of the VUnc (see Fig 3.1 third panel). DGgcl is defined as the voxels that border mVUnc as well as the clear label medial to the hippocampus. Along the inferior side of the uncus, DGgcl within the vertical component of the uncus will overlap with DGgcl in the rest of the hippocampus. Thus, the two labels are continuous after being combined.




Fig 3.3. Borders of mVUnc. This is a dummy label defined as everything medial to the anterior two thirds of the vertical component of the uncus (approximated by the blue line in the left panel).

Indusium griseum. This is the grey matter that extends from the posterior and medial edge of the hippocampal tail. This part may have already be labeled as part of this as hippocampal grey matter after dilating the SRLM label, but if not simply label all grey matter in this vicinity as indusium griseum, and the border between indusium griseum and hippocampus will be established based on the 3D model in the following way: viewing the hippocampus from the superior side, follow the curvature of the posterior edge of the hippocampal tail. When the tail begins to curve posteriorly again, extend the medial curve to cut it off, forming indusium griseum to the posterior of that line. In the current data, Indusium griseum was labelled only where it immediately borders hippocampus.




Fig 3.4. Indusium griseum border. Top row shows grey matter extension from the posterior hippocampus. Bottom row shows the border defined from a superior view of the hippocampal tail (white line)

Appendix 2 (Chapter 2): Manual subfield segmentation by alignment with Ding & Van Hoesen, 2015; Duvernoy *et al.*, 2013

Most of this work was completed before the release of segmentation protocols that make use of recent histological work by Ding *et al.* (namely Dalton *et al.*, 2017 and Berron *et al.*, 2017), and so we had no clear protocol available that could leverage this new information. We aimed to capture the highest subfield validity, in this case at the expense of quantified reliability, so we thus performed our segmentation by trying to find the best correspondences between slices in our data and in Ding & Van Hoesen's images.

Segmentations were performed by raters JD and KF. Each coronal slice from the hippocampal head was matched to its closest corresponding slice from the hippocampus with the same number of anterior digitations in Ding & Van Hoesen. Some Coronal slices from the hippocampal body were matched to Ding & Van Hoesen or Duvernoy *et al.*, depending on which image most closely resembled that participant. In the hippocampal tail, slices were segmented in the sagittal view and matched to descriptions and images found in Duvernoy *et al.* After slices were matched between our images and reference images, the segmenter subjectively interpolated the subfield borders over the surface of the 3D model in order to ensure the subfields followed a smooth trajectory. In some cases, borders were roughly matched to differences in grey matter intensity, with areas CA2, CA3, and the subiculum sometimes appearing slightly darker than CA1 and the dentate gyrus. However, these differences were very inconsistently visible and so were only used when they closely corresponded to expected border locations as per histological reference.

Note that not every slice in our data had a perfectly corresponding image in the anatomical references, in part due to the lack of fully 3D high resolution histological references. Also, in some cases there were still considerable differences between slices of the current dataset and the histological references. This might be due to any number of factors including inter-individual differences, tissue changes during histological processing, or different slice angles. In these cases, the subjective opinion of the rater was used, in conjunction with reference to where the subfield surfaces might be expected to fall on the surface of the 3D model.




Fig 1. Example slices showing corresponding subfield arrangements between anatomical references (Ding & Van Hoesen, 2015 and Duvernoy *et al.*, 2013) and segmentations performed on our dataset (same example hippocampus as shown throughout). Rows 1-7 are coronal and row 7 is sagittal, reflecting the plane in which subfield segmentation was performed. Left column shows examples of anatomical references (Ding & Van Hoesen for rows 1-5, Duvernoy *et al.* for rows 6-7), middle column shows our dataset

with the grey matter (red), SRLM (green), and cysts (yellow), and right column shows example segmentations by rater JD with subiculum in brown, CA1 in blue, CA2 in white, CA3 in pink, and dentate gyrus in cyan.

Appendix 3 (Chapter 2): Manual histology labelling protocol

Slices were stained with NeuN, and delineation of the hippocampal subfields was performed according to the rules set in Ding & Van Hoesen (2015). The border of the subiculum with the entorhinal cortex is shown by the smaller neurons in the subiculum compared with the entorhinal cortex (Fig 1).




Fig 1. Left: histology slice from the hippocampal head stained with NeuN. Right: zoomed in to the border between the entorhinal cortex and subiculum. Green and red arrows indicate larger and smaller layer 2 neurons in entorhinal cortex and subiculum, respectively.

The border of the subiculum with CA1 can be determined because the subiculum has a polymorphic cell layer in addition to a pyramidal cell layer, which is a distinct characteristic from CA1 (Fig 2).




Fig 2. Left: histology slice from the hippocampal head including the posterior uncus, stained with NeuN. Right: zoomed in to the border between subiculum and CA1

The tissue from this particular sample shows marked cellular loss in CA1 and, to a lesser extent, CA3, but sparing CA2 (Fig 3 left) which is typical of hippocampal epilepsy pathology (Blümcke *et al.*, 2007). In addition to the difference in cellular loss between

CA2 and CA3, there is also a shallow depression in the pyramidal cell layer between CA2 and CA3 to help determine this border (Fig 3 middle).




Fig 3. Left: view of CA1, CA2, CA3, DG/CA3h subfields with arrows indicating cellular loss in CA1 and, to a lesser extent, CA3. Middle: zoomed in to the border between CA2 and CA3. Right: Adapted from Ding & Van Hoesen (2015). Arrows in the middle and right panels indicate the shallow depression in pyramidal cell layer between CA2 and CA3

CA3 has darker and more densely packed neurons than CA3h region (Ding & Van Hoesen, 2015). This was less evident given the cell loss in CA3, but some difference was still visible (Fig 4 left). The rest of the DG/CA3h border is determined by the dense granule cell layer of the DG (Fig 4 right).




Fig 4. Left: border between CA3 and CA3h. Right: entire DG/CA3h subfield including the granule cell layer of the DG.

This paper also considers the modified subfields that Ding & Van Hoesen (2015) discuss in the head of the hippocampus. For example, CA3(uncal) and DG/Ca3h(uncal) are very similar to their counterparts in the rest of the hippocampus, but the cells are slightly less densely packed and slightly less darkly stained with NeuN (Fig 5).




Fig 5. The uncus of the hippocampus. The uncal modified subfields are outlined: CA3 in blue, DG/CA3h in pink. Note that the subfields are similar but less dense to those in the rest of the hippocampus (i.e. Figures 1 to 4).

Appendix 4 (Chapter 3): Additional anatomical details of the hippocampus and surrounding structures in BigBrain

The dentate gyrus, modelled in Fig 1A, was excluded from our hippocampal unfolding framework because it was used as a boundary condition for proximal-distal coordinates. However, it was easily differentiated from other subfields by its very high neuronal density. Fig 1B shows residual staining at tissue boundaries, as well as some staining within the SRLM laminae which may be due to the presence of inhibitory interneurons.




Fig 1. Anatomical details noted in BigBrain 3D histology outside of ‘archicortex’ label. A) 3D models of the dentate gyrus from the superior (top) and inferior (bottom), with the red line demonstrating the ‘U’ shape and green arrows indicating ‘dentations’ that are prominent in the human DG and follow the gyrification of archicortex. B) Increased staining can be seen on CSF boundaries (e.g. orange arrows) but also may include interneurons found in the stratum moleculare of the CA fields (blue arrows) and stratum moleculare of the dentate gyrus (green arrows), which are intermittently fused or separated by the vestigial hippocampal sulcus. Images shown are from the original 20um resolution slices (left) and a coronal slice from the 40um isoluminant hippocampal block (right).

We detected gyral peaks in the left and right hippocampi and calculated the distance between them. Briefly, peaks were detected by taking an anterior-posterior profile midway through manually defined subfield CA1 (where most gyri were centered) and detecting local maxima in curvature (see sections 2.3 and 2.4). Distances between gyral peaks are shown in Fig 2, and were generally low in the uncus, high in the remaining hippocampal head and anterior body, and decreasing in size through the body and tail of the hippocampus. Smaller gyri had as little as 0.4mm of tissue (alveus on the outer surface or SRLM on the inner surface) separating them. Thus, it should be

possible to visualize all these gyri with 0.4mm isotropic resolution, provided sufficient image contrast is available. However, with more sophisticated modelling it may be possible to estimate these gyri despite partial voluming between gyri with resolutions as low as the Nyquist frequency between gyri ($\sim 1\text{mm}$ isotropic).





Fig 2. Peak-peak distances between hippocampal gyri.

Fig 3 shows the distributions of all these data from both the left and right hippocampi, colour-coded by subfield. Linear trends where R-squared values were greater than 0.1 are overlaid. Laminar feature Mean(y), or the mean amount of staining, increased towards the posterior of the hippocampus in regions CA1 and CA4. Mean(x.d), or the depth at which the greatest change in staining intensity was seen, increased in CA2 and CA3 towards the posterior as well. Skew(x), or the skew on the depth of neurons, decreased towards posterior subiculum. The morphological feature gyrification also showed a decrease towards the posterior in most prominently in CA1 which was also noted in Fig 1, but also in CA3.



<2 column width> Fig 3. Examination of anterior-posterior differences in hippocampal structure. The x-axes represent the anterior-posterior axis of the hippocampus and the y-axes are the corresponding feature values, colour-coded by manual subfields. Linear correlations with R-squared values greater than 0.1 are overlaid.

Appendix 5 (Chapter 3): Additional methodological details

Our Laplacian unfolding framework was modified in several ways since its use in our previous publication (DeKraker *et al.* 2018):

- 1) DG was excluded from the domain of hippocampal archicortex (though CA4 was instead differentiated from CA3).
- 2) Unfolded representations were modified to have a 2:1 rather than 1:1 anterior-posterior:proximal-distal aspect ratio which we found to more closely match the real world dimensions of hippocampal tissue. BigBrain flatmaps were generated with 512x256 coordinate points in order to capture the higher level of detail available in that dataset (previous work was 100x100 points).
- 3) General code optimizations to make it faster given large datasets and more robust to errors in manual segmentation. Interpolated mid-surfaces are now also automatically generated. These changes do not affect the unfolding outputs.

Updated code and the history of changes made can be seen at <https://github.com/jordandekraker/HippUnfolding>.

The Equivolume model solution was obtained by a python wrapper of ‘CBStools’ called ‘Nighres’ (<https://github.com/nighres/nighres>).

Laminar sampling differences between Wagstyl *et al.* (2018) and the current study differed in the following way: horizontal (i.e. perpendicular to each profile) smoothing was performed with a Gaussian kernel of $\sigma=3$ in unfolded space, rather than with a parametrically optimized σ in native space (σ could not be optimized to detect 3-5 laminar profile peaks in the neocortex since we expected only 1-3 laminar profile peaks in the archicortex and we wanted to avoid overfitting).

Mean curvature was estimated using the function ‘patchcurvature()’ from the MathWorks File Exchange (<https://www.mathworks.com/matlabcentral/fileexchange/32573-patch-curvature>).

Prior to smoothing of features in unfolded space, the unfolded space was reparametrized to reflect real-world distances between points, ensuring that all regions of the hippocampus were smoothed with equal kernel sizes. This was done by scaling the distances between points in unfolded space by their distances in native space. This is illustrated in Fig 4. Smoothed flatmaps were then returned to standard unfolded space.





Fig 4. Illustration of one feature (thickness) in unfolded space (left) in the reparametrized unfolded space (right). Note that the axes in the reparametrized unfolded space reflect real-world distances (*0.04mm).

Appendix 6 (Chapter 3): Laminar modelling using Laplacian vs. Equivolume model in BigBrain

The Equivolume model aims to account for displacements of laminae due to curvature. The Laplacian model remains smooth even over complex shapes but does not account for displacement of laminae due to curvature, and so there is an offset between the depths at which peak staining along a laminar profile is found. The Equivolume model more closely aligns peaks between gyral and sulcal areas but is highly susceptible to distortion from the rough surfaces of detailed manual segmentations (e.g. due to subicular 'islands'), which can be seen in the sagittal slice of the model solution.




Fig 5. Comparison of Equivolume and Laplacian models of archicortical lamination. The left images show a sagittal slice of each model solution in the right hemisphere. Right images show two example laminar profiles, one at the peak of a gyrus (A-P,P-D coordinates 236,104) and one at the depth of sulcus (A-P,P-D coordinates 218,104) in the CA1 of the right hippocampus.

Appendix 7 (Chapter 3): Additional clustering results

PCA followed by K-means clustering is only one of many possible methods for extracting summary information from the many features examined in the current study. In this section, we explore other possible methods

Fig 6 demonstrates the discriminability between gold standard subfield definitions using the first 8 PCA components. For example, comparing PCA component 1 against any other component, we can see that CA2 shows high values while subiculum shows low values. In combination with other PCA components, linear decision boundaries can be drawn between the subfields according to their scores along each component. This linearity makes clustering a well posed problem, suitable for k-means or other similar algorithms.

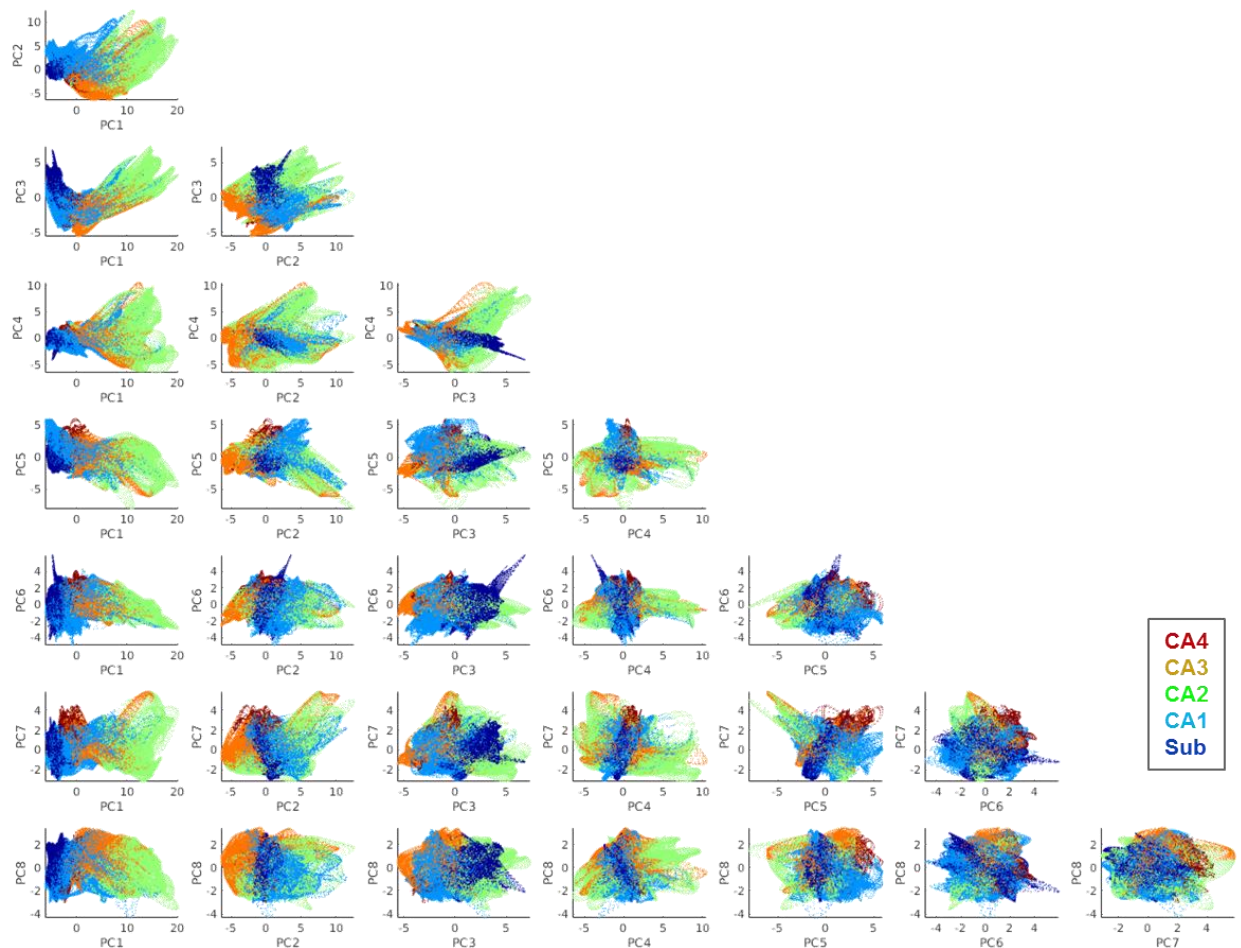
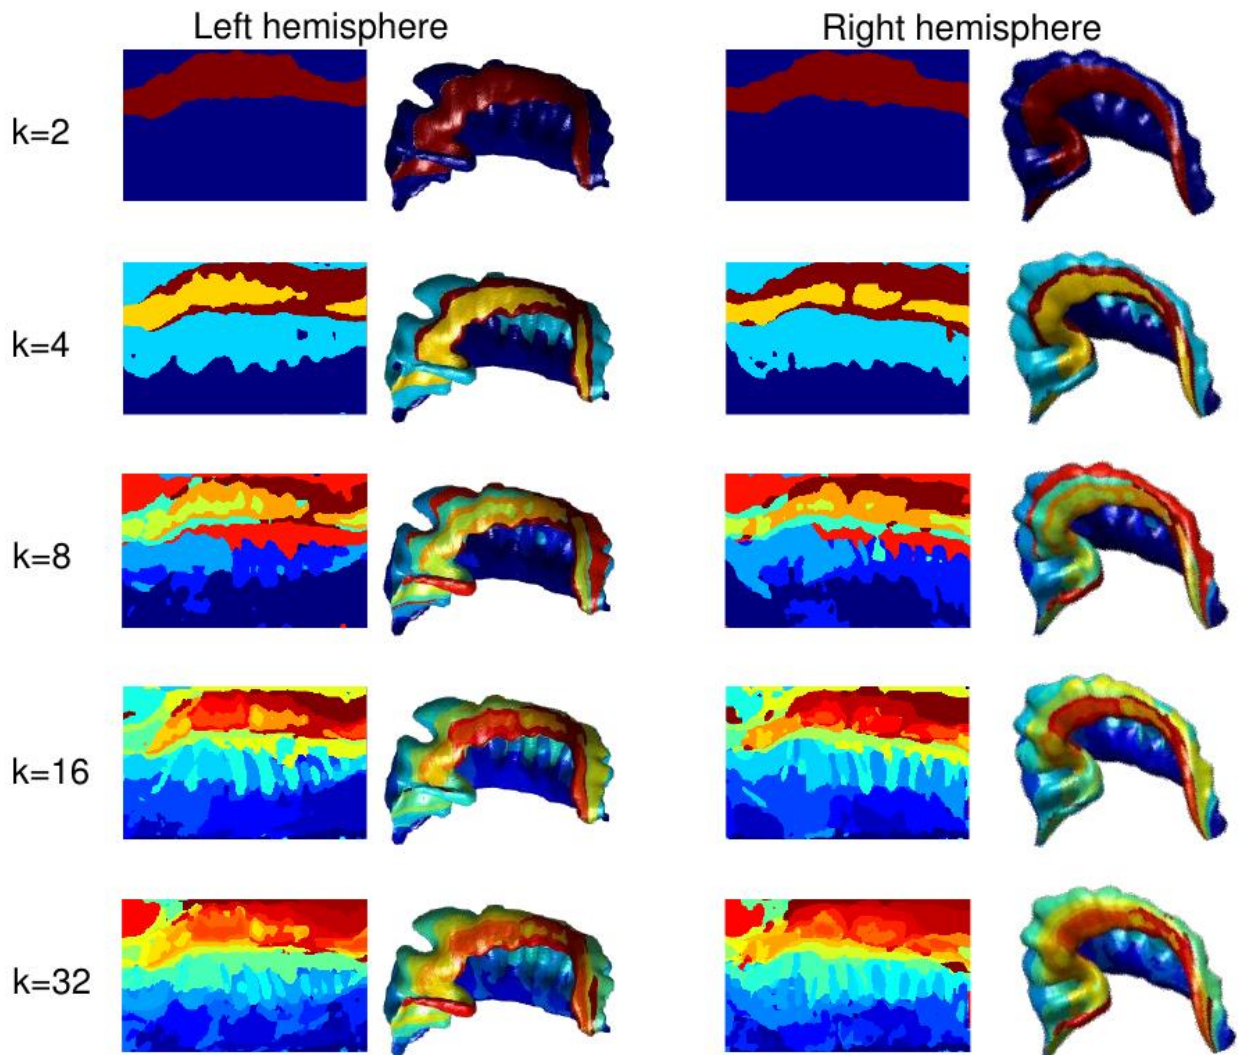


Fig 6. Linear separability of subfields using only PCA components. The first 7 PCA components are each plotted against each other, and with points coloured according to the gold standard subfield definitions. Blue indicates subiculum; cyan indicates CA1; green indicates CA2, orange indicates CA3, and red indicates CA4.

By setting $k=5$ in k-means clustering we imposed some prior information onto the segmentation of the hippocampus. Thus, we also explored different possible parcellation schemes by setting $k=[2,4,8,16,32]$. We ordered the resulting clusters according to their median proximal-distal distance in order to better align with our manual segmentation label scheme. Within each clustering result, borders can be seen near the manually defined subfield borders. With greater numbers of clusters, the subfields divide into additional regions, but most borders separate proximal-distal regions rather than anterior-posterior regions. This could be due to the presence of transition zones between subfields, or perhaps in the case of subiculum additional proximal-distal segments become clustered due to subicular subregions (e.g. prosubiculum, presubiculum, subiculum proper, or parasubiculum). Overall, these data provide strong motivation for the segmentation of the hippocampus along its proximal-distal rather than its anterior-posterior axis, and the most seen borders resemble the classic hippocampal subfield definitions. Plotting Dice overlap as a function of the number of clusters shows an elbow for most subfields near $k=4$, which suggests that our *a-priori* choice of $k=5$ is likely an appropriate number of clusters.






Fig 7. Exploration of different clustering schemes by varying k in k -means clustering. Resulting clusters are labeled and coloured according to their median proximal-distal distance. Below are plotted Dice scores for each subfield by the number of clusters, k , in k -means clustering.

In addition to k -means, we also tested whether another clustering algorithm might yield similar results, hierarchical clustering. Distances were calculated on the first 8 PCA components, using a 'ward' distance metric. Because of the large size of the data, each PCA component was reshaped into a flatmap (as in Fig 5 in the main text) and then downsampled by a factor of half before calculating distances. Fig 8 (left) shows possible clustering hierarchical solutions based on these distances. Choosing 5 clusters, we see a similar pattern of clustering as when K -means was used (Fig 8 right).




Fig 8. Hierarchical clustering of PCA component 1 to 8. The left image shows a dendrogram of the linkage between the first 30 possible cluster solutions, with the dotted line indicating a divide into 5 clusters (colour coded according to the clusters shown on the right). The right shows a clustering of this linkage into 5 clusters, projected onto unfolded and a mid-surface space.

Data-driven clustering results were evaluated using Dice overlap in unfolded space, however, this does not account for different thicknesses of the different subfields

or the fact that clusters had to be extrapolated over the anterior and posterior 10% edges which were not analyzed due to high noise. We thus additionally projected clusters to native space using nearest neighbour interpolation and calculated Dice overlap scores again, as shown in Table 1. Note that these scores are generally lower than those reported in the main body of the paper, which is most likely due to the fact that subfields which were most poorly captured by clustering, CA2 and CA3, are thinnest and therefore small deviations have a larger impact on overall volume overlaps. Furthermore, extrapolation at the anterior and posterior edges of the hippocampus likely contributed to reduced overlap, since clustering results were not particularly close to manual gold standard near these edges, perhaps due to high noise. These regions also contained the highest distortion in unfolded space (e.g. Appendix 4 Fig 3), which may have also contributed to the difference in Dice scores between native and unfolded space.

Table 1. Dice overlap scores between k-means clusters using all features and their closest corresponding manually defined subfield.

	Left	Right
Sub	0.76	0.77
CA1	0.67	0.65
CA2	0.65	0.64
CA3	0.38	0.32
CA2&CA3	0.61	0.51
CA4	0	0

Appendix 8 (Chapter 4): Acquisition and preprocessing details of datasets used in U-Net training

All datasets examined employ T2w sub millimetric MRI data. Acquisition and preprocessing details were as follows:

UWO7T:

Four T2-weighted turbo spin echo (TSE) 3D (3D sagittal, matrix: 260×366 , 266 slices, 0.6mm^3 isotropic, ~ 8.5 mins per scan) images were acquired from each participant. All images were acquired in sagittal rather than coronal oblique orientation for optimal whole brain coverage. The first T2-weighted image (scan 1) was upsampled to 0.3mm^3 isovoxels using cubic spline interpolation; subsequently, scans 2, 3, and 4 were rigidly registered to scan 1 using FSL FLIRT registration (Jenkinson, 2002, Tofts, 2005). All four scans were then averaged together to produce a single, 0.3mm^3 isovoxel, high-contrast volume.

UWO3T:

T2-weighted SPACE (3D sagittal, 0.8mm^3 isotropic, iPAT GRAPPA acceleration x2) images were acquired from each participant. MRI data were converted to brain imaging data structure (BIDS) and run through fmriprep-v1.1.8 (Esteban et al., 2018).

UPennExVivo:

31 specimens were collected from 25 donors. The donors were older adults with and without dementia (nine AD, nine other dementia, and seven no dementia). Specimens, fixed in 10% formalin solution for a minimum of 21 d, were imaged on a Varian 9.4-Tesla animal scanner at $200 \times 200 \times 200 \mu\text{m}^3$ ($160 \times 160 \times 160 \mu\text{m}^3$ in one specimen) resolution. Only 22 samples were used in the present work since other samples were missing some hippocampal tissue which could lead to distortions in our unfolding coordinate framework. See Adler *et al.*, (2018) for further details.

HCP1200:

T2w image were acquired using the variable flip angle turbo spin-echo sequence (Siemens SPACE (Mugler *et al.*, 2000)) with 0.7mm isotropic resolution (same matrix, FOV, and slices as in the T1w), $\text{TR}=3200\text{ms}$, $\text{TE}=565\text{ms}$, $\text{BW}=744$ Hz per pixel, no fat suppression pulse, phase encoding undersampling factor GRAPPA=2, total turbo factor=314 (to be achieved with a combination of turbo factor and slice turbo factor, when available), echo train length of 1105 echoes, 10% phase encoding oversampling (anterior - posterior) to avoid nose wrap-around, readout along superior - inferior direction with dwell time of $2.1\mu\text{s}$ (for readout distortion correction with FUGUE). See Glasser *et al.*, (2013) for further details.

Curriculum Vitae

Jordan DeKraker

Academic Degrees

- Currently completing PhD in Neuroscience at Western University (2016-present)
- MSc in Behavioural and Cognitive Neuroscience at Western University (2016)
- BSc in Neuroscience Combined Honours at Carleton University (2014)
- French Immersion high school diploma at Parkside Collegiate Institute (2010)

Awards and Scholarships

- Post-Graduate Scholarship - PhD (NSERC) in 2016 - 2019 (Western University)
 - Valued at \$105 000 over three years
- Doctoral Excellence Research (Western University) in 2016-2019
 - Valued at \$20 000 over three years
- Ontario Graduate Scholarship in 2015 (Western University)
 - Valued at \$15 000 over one year
- Canadian Graduate Scholarship – MSc (NSERC) in 2014 (Western University)
 - Valued at \$17 500 over one year
- Undergraduate Student Research Awards (NSERC) in 2012 (Carleton University); and again in 2013 (Western University).
 - Valued at \$4,500 from NSERC + minimum \$1,125 from the supervisor over one semester
- Dean's Summer Research Internship (Carleton University) in 2011
 - Valued at \$4620 over one semester
- Carleton Academic Scholarship for high school grades 85-89.9% and for maintaining a minimum GPA of 10.0 (2010 - 2014)
 - Valued at \$8,000 over four years
- Queen Elizabeth II Aiming for the Top Tuition Scholarship (OSAP), based on financial need and academic excellence (2010 - 2014)

Publications

- Paquola, C., Benkarim, O., **DeKraker, J.**, Lariviere, S., Frässle, S., Royer, J., Valk, S., Bernasconi, A., Bernasconi, N., Khan, A., Razi, A., Smallwood, J., Bernhardt, B. (in prep). Distinct dimensions of cortical function converge along cytoarchitectural and long-axis gradients of the mesiotemporal lobe.
- Hussain, U., **DeKraker, J.**, Rajakumar, N., Baron, C.A., Khan, A.R. (bioRxiv). Diffusion MRI of the Unfolded Hippocampus. (preprint April 7, 2020)
- **DeKraker, J.**, Lau, J. C., Ferko, K. M., Khan, A. R., Köhler, S. (2020). Hippocampal subfields revealed through unfolding and unsupervised clustering

of laminar and morphological features in 3D BigBrain. *Neuroimage* 206 (2020): 116328.

- Olsen, R. K., *et al.* (2019). Progress Update from the Hippocampal Subfields Group. *Alzheimer's & Dementia: Diagnosis, Assessment & Disease Monitoring*.
- Xiao, Y., Lau, J.C., Anderson, T., **DeKraker, J.**, Louis Collins, D., Peters, T.M., and Khan, A.R. "An accurate registration of the BigBrain dataset with the MNI PD25 and ICBM152 atlases," *Scientific Data*, 6(1), 1-9, 2019.
- **DeKraker, J.**, Ferko, K. M., Lau, J. C., Köhler, S., Khan, A. R. (2018). Unfolding the hippocampus: An intrinsic coordinate system for subfield segmentations and quantitative mapping. *NeuroImage*, 167, 408-418.
- Fiacconi, C. M., **DeKraker, J.**, Köhler, S. (2015). Psychophysiological evidence for the role of emotion in adaptive memory. *Journal of Experimental Psychology: General*, 144(5), 925.
- Davis-MacNevin, P. L., **DeKraker, J.**, LaDouceur, L., Holahan, M. R. (2013). Comparison of the MK-801-induced increase in non-rewarded appetitive responding with dopamine agonists and locomotor activity in rats. *Journal of Psychopharmacology*, 0269881113492029.
- Mentioned in *Brain Research* for contributions to the study "Enhanced adolescent learning and hippocampal axonal projections following preadolescent spatial exposure to a water or dry maze" (Wartman *et al.*, 2012).

Oral presentations

- "Computational unfolding of the human hippocampus to derive subfields in MRI" at:
 - Helmholtz International BigBrain Analytics Learning Laboratory, Montreal, 2020 (international) (invited but not yet delivered)
 - Canada Student Health Research Forum, Winnipeg, 2020 (national) (invited but not yet delivered)
- "Computational anatomy of the hippocampus: bridging spatial scales with topological (archi)cortical modelling" at:
 - National Institute of Mental Health, Washington DC, 2019 (invited talk by Dr. Armin Raznahan)
- "Intro to Deep Learning" (theoretical component, assisted in practical component) at:
 - Brainhack Western, London, 2019 (regional)
- "Computational Anatomy of the 3D BigBrain Hippocampus" at:
 - Hippocampal Subfields Group meeting, Magdeburg, 2018 (international)
- "Hippocampal unfolding and histological characterization using data-driven methods developed for the neocortex" at:
 - BigBrain workshop, Montreal, 2017 (international)
- "Unravelling subfields of the hippocampal head with 7T MRI: leveraging the dark band" at:
 - Society for Neuroscience nanosymposium, San Diego, 2016 (international)

- Hippocampal Subfields Group meeting, San Diego, 2016 (international)
- “Unravelling Subfields of the Hippocampal Head with 7T MRI” at:
 - Epilepsy Research rounds, Robarts Research Institute, 2016 (Western University)
 - Coffee talk series, Brain and Mind Institute, 2016 (Western University)
- “Imaging the human hippocampal head with 7 Tesla MRI at ultra-high resolution” at:
 - Epilepsy and Neuroscience Research Day, London, ON, Canada 2015 (Regional)

Poster presentations

- “Automated hippocampal unfolding for quantitative mapping, morphometry, and subfield definition” at:
 - Organization for Human Brain Mapping, Montreal, 2020 (international)
 - CIHR National Health Research Poster Presentation, Manitoba, 2020 (national)
- “Morphological and laminar characterization of the human hippocampus in 3D BigBrain” at:
 - Organization for Human Brain Mapping, Rome, 2019 (international)
 - Canadian Association for Neuroscience, Toronto, 2019 (national)
- “Meso-scale anatomy of the hippocampus in BigBrain 3D histology: comparison to ultra-high resolution MRI and data-driven segmentation” at:
 - Hippocampal Subfields Group meeting, Magdeburg, 2018 (international)
- “Using Laplace’s equation to generate a standardized, ‘unfolded’ map of hippocampal grey matter” at:
 - Organization for Human Brain Mapping, Vancouver, 2017 (international)
- “Unfolded Hippocampal coordinate system for Quantitative mapping and subfield segmentation” at:
 - Hippocampal Subfields Group meeting, San Diego, 2016 (international)
- “Unravelling Hippocampal Subfields with 7T MRI” at:
 - London Imaging Discovery Day in London, 2015 (Regional)
- “Autonomic Contributions to Adaptive Memory” at:
 - Cognitive Neuroscience Society in San Francisco, 2015 (international)
 - Canadian Society for Brain, Behaviour and Cognitive Sciences in Ottawa, 2015 (National)
 - Lake Ontario Visionary Establishment in Niagara Falls, 2015 (regional)
 - Inaugural Brain and Mind Symposium in London, 2015 (Western University)

Co-authored poster presentations (highlights only)

- “Piriform Cortex Parcellation through Unfolding and Clustering of Laminar Features in 3D BigBrain” (first author Nickolas Christidis) at:
 - Organization for Human Brain Mapping, Montreal, 2020 (international)

- “Cortical confluence: Cytoarchitectural mapping of the transition from iso to allocortex” (first author Casey Parquola) at:
 - Organization for Human Brain Mapping, Montreal, 2020 (international)
- “Characterization of the hippocampal formation using diffusion-weighted imaging” (first author Mohamed Yousif) at:
 - Organization for Human Brain Mapping, Montreal, 2020 (international)

Teaching

- Co-supervised co-op students:
 - Bilal Amin, University of Waterloo: Classification and Segmentation of Focal Cortical Dysplasias in Histology using Deep Convolutional Neural Networks (2020)
- Tutorial series presentations:
 - “Image segmentation with few training samples using model-based augmentation” tutorial talk and interactive session (Western University and London Health Sciences Center, May 2019)
- Co-supervised BSc Honours projects:
 - Sanveer Dhanju, Western University: Segmentation of the Hippocampus in MRI using Convolutional Neural Networks (2018)
 - Kayla Ferko, Western University: Segmentation of Hippocampal Subfields Using 3-Tesla MRI (2016)
 - Sawayra Owais, Western University: Psychophysiological Correlates of Adaptive Memory (2015)
- Teaching assistant:
 - Psychology of Perception (PSYCHOL 2015B) 2016, Western University
 - Evolutionary Psychology (PSYCHOL 3229A) 2014, Western University
 - Introduction to Psychology (PSYCHOL 1000) 2014, Western University

Leadership

- Winner of the Open Science Award at Brainhack Western, 2019. Project title: “Artificially Intelligent Autobiographical Interview”
- Organizer and project leader at Brainhack Western, 2018. Project title: “Automated single-trial modelling of fMRI data in BIDS format”
- Brainhack Vancouver participant, 2017
- Vice President external and co-founder of the Carleton Neuroscience Society (CNS) 2012-2014
- Member of the Carleton University Varsity Fencing Team, 2011-2013
- Winner of Athletic Bursary from Carleton in 2012; 2013
- Tutor in Introductory Chemistry at Carleton University, 2011-2013
- Summer camp counselor and seasonal labourer at Farmtown Canada, 2007-2010

Co-authored presentations & unpublished work

- Christidis, N.K, **DeKraker, J.**, Xiao, Y, *et al.* “Piriform Cortex Parcellation through Unfolding and Clustering of Laminar Features in the 3D BigBrain” Poster presentation at Organization for Human Brain Mapping, Montreal, Canada, 2020 (international).
- Wang, B. T., **DeKraker, J.**, Khan, A. R. “A quantitative comparison of registration techniques using an intrinsic hippocampal coordinate space” (prepped for publication in PloS One in 2018, but the project was dropped as the first author left for med school)