

Accurate Bayesian segmentation of thalamic nuclei using diffusion MRI and an improved histological atlas

Henry F. J. Tregidgo^{a,*}, Sonja Soskic^a, Juri Althonayan^a, Chiara Maffei^b, Koen Van Leemput^{b,c}, Polina Golland^d, Anastasia Yendiki^b, Daniel C. Alexander^e, Martina Bocchetta^{f,g}, Jonathan D. Rohrer^f, Juan Eugenio Iglesias^{a,b,d}, for the Alzheimer's Disease Neuroimaging Initiative¹

^aCentre for Medical Image Computing, Department of Medical Physics and Biomedical Engineering, University College London, UK,

^bMartinos Center for Biomedical Imaging, Massachusetts General Hospital and Harvard Medical School, USA,

^cDepartment of Applied Mathematics and Computer Science, Technical University of Denmark, Denmark,

^dComputer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, USA,

^eCentre for Medical Image Computing, Department of Computer Science, University College London, UK,

^fDementia Research Centre, Department of Neurodegenerative Disease, UCL Queen Square Institute of Neurology, University College London, UK,

^gCentre for Cognitive Neuroscience, Department of Life Sciences, College of Health, Medicine and Life Sciences, Brunel University London, UK,

ARTICLE INFO

Keywords:

Thalamus
Atlasing
Diffusion MRI
Segmentation
Bayesian inference

ABSTRACT

The human thalamus is a highly connected brain structure, which is key for the control of numerous functions and is involved in several neurological disorders. Recently, neuroimaging studies have increasingly focused on the volume and connectivity of the specific nuclei comprising this structure, rather than looking at the thalamus as a whole. However, accurate identification of cytoarchitectonically designed histological nuclei on standard *in vivo* structural MRI is hampered by the lack of image contrast that can be used to distinguish nuclei from each other and from surrounding white matter tracts. While diffusion MRI may offer such contrast, it has lower resolution and lacks some boundaries visible in structural imaging. In this work, we present a Bayesian segmentation algorithm for the thalamus. This algorithm combines prior information from a probabilistic atlas with likelihood models for both structural and diffusion MRI, allowing label boundaries to be informed by both modalities. We present an improved probabilistic atlas, incorporating 26 thalamic nuclei identified from histology and 45 white matter tracts identified in ultra-high gradient strength diffusion imaging. We present a family of likelihood models for diffusion tensor imaging, ensuring compatibility with the vast majority of neuroimaging datasets that include diffusion MRI data. The use of these diffusion likelihood models greatly improves identification of nuclei versus segmentation based solely on structural MRI. Dice comparison of 5 manually identifiable groups of nuclei to ground truth segmentations show improvements of up to 10 percentage points. Additionally, our chosen model shows a high degree of reliability, with median test-retest Dice scores above 0.85 for four out of five nuclei groups, whilst also offering improved detection of differential thalamic involvement in Alzheimer's disease (AUROC 83.36%). The probabilistic atlas and segmentation tool will be made publicly available as part of the neuroimaging package FreeSurfer.

1. Introduction

2 The thalamus has traditionally been considered a relay
3 station for information in the brain, with extensive connections
4 to both cortical and subcortical structures (Schmahmann,
5 2003; Behrens et al., 2003). As such, it integrates
6 information processing between cortical regions (Sherman,
7 2007, 2016; Hwang et al., 2017) and is associated with
8 a wide range of functions including cognition, memory,
9 sensory and motor functions, regulation of consciousness
10 and spoken language among others (Sherman and Guillery,
11 2001; Schmahmann, 2003; Fama and Sullivan, 2015). Ad-
12 ditionally, neurodegenerative pathological processes in the

1 thalamus have been associated with Alzheimer's disease
2 (AD) (de Jong et al., 2008; Zarei et al., 2010), frontotemporal
3 dementia (Bocchetta et al., 2018; McKenna et al., 2022),
4 Huntington's disease (Aron et al., 2003; Kassubek et al.,
5 2005) and multiple sclerosis (Minagar et al., 2013; Planche
6 et al., 2019).

7 With such wide established connections and functions,
8 the thalamus is a frequent target in MRI-based neuroimaging
9 studies and a focus for research in relation to both healthy
10 and disordered brain function. This creates a need for reliable
11 identification of thalamic borders. Therefore, the thalamus
12 is defined by several structural MRI (*sMRI*) segmentation
13 methods, including multi-atlas segmentation (Heckemann
14 et al., 2006), Bayesian segmentation (Puonti et al., 2016)
15 and convolutional neural networks (*CNNs*) (Wachinger et al.,
16 2018; Billot et al., 2020; Henschel et al., 2020). Additionally,
17 the thalamus has been included in popular image processing
18 packages, including FreeSurfer's (Fischl, 2012) recon-all
19 stream, which uses a probabilistic atlas of anatomy and MRI
20 intensity (Fischl et al., 2002), and the FMRIB Software
21 Library (*FSL*) (Smith et al., 2004), which includes a model

 h.tregidgo@ucl.ac.uk (H.F.J. Tregidgo)

ORCID(s): 0000-0002-3509-8154 (H.F.J. Tregidgo);

0000-0003-1814-5024 (M. Bocchetta); 0000-0002-6155-8417 (J.D. Rohrer)

¹Data used in this article are partly from the Alzheimer's Disease Neuroimaging Initiative database (<http://adni.loni.usc.edu>). Investigators in the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis of this report. A complete listing of investigators can be found at: adni.loni.usc.edu/wp-content/ADNI_Acknowledgement_List.pdf

Multi-modal thalamic segmentation

1 of shape and appearance in its implementation (FIRST)
2 (Pateraude et al., 2011).

3 The methods above segment the thalamus as a single
4 label, however in reality it is a complex and heterogeneous
5 structure. It is composed of 14 major nuclei, which may
6 be split further into 50 subnuclei depending on the level of
7 detail in the classification and agreement on neuroanatomical
8 definition. There are multiple such definitions with
9 varying numbers of subnuclei (Morel, 2007; Jones, 2012;
10 Mai and Majtanik, 2019). These nuclei have distinct pat-
11 terns of connections with other brain regions and subserve
12 different functions, including associative, sensory, motor,
13 cognitive and limbic (Schmahmann, 2003). For example,
14 the ventral lateral posterior nucleus is involved in motor
15 function through connections with the cerebellum and the
16 motor cortex, while the mediodorsal nucleus has connec-
17 tions with the prefrontal cortex and plays a role in cognitive
18 and emotional processes (Mai and Forutan, 2012; Schmahn-
19 mann, 2003). In addition, neuropathological studies have
20 demonstrated preferential involvement of certain thalamic
21 nuclei in several conditions, such as the caudal intralaminar
22 nuclei in Parkinson's disease (Henderson et al., 2000), the
23 anterior nuclei in AD (Braak and Braak, 1991a,b), and the
24 pulvinar in the *C9orf72* genetic subtype of frontotemporal
25 dementia (Vatsavayai et al., 2016). These studies provide
26 strong motivation for the design of automated segmentation
27 algorithms that accurately define thalamic nuclei *in vivo*,
28 enabling identification of reliable and precise biomarkers.

29 Different approaches have been used to segment thala-
30 mic nuclei. There are segmentation strategies that attempt
31 to directly register histology derived labels to MRI. For
32 instance, manually labelled histology can be used to generate
33 a reference space atlas that may then be applied to *in vivo*
34 MRI through registration-based segmentation (Krauth et al.,
35 2010; Jakab et al., 2012; Sadikot et al., 2011). However,
36 such approaches are limited by the difficulty in registering
37 MR images with different contrasts. Other techniques define
38 their label scheme based on information derived from the
39 imaging data to be segmented. For example, diffusion MRI
40 (*dMRI*) has been used to define thalamic regions by cluster-
41 ing voxels based on diffusion tensor imaging (*DTI*) indices
42 (Mang et al., 2012) and orientation distribution functions
43 (Battistella et al., 2017; Semedo et al., 2018). Other studies
44 have divided the thalamus into regions based on their
45 cortical connectivity, either through resting-state functional
46 MRI time course correlations (Zhang et al., 2008) or *dMRI*
47 tractography (Behrens et al., 2003; Johansen-Berg et al.,
48 2005). However, exactly how thalamic regions defined by
49 functional MRI relate to neurobiology is not fully under-
50 stood (Eickhoff et al., 2015) and there is some indication
51 that tractography-based segmentations are insensitive to the
52 internal structure of the thalamus (Clayden et al., 2019).

53 The development of advanced MRI acquisitions has also
54 allowed for atlases to be defined from manual segmentation
55 of *in vivo* imaging directly, due to improved resolution and
56 contrast. For example, guided by histological atlases, it has
57 been possible to manually identify nuclei on advanced sMRI

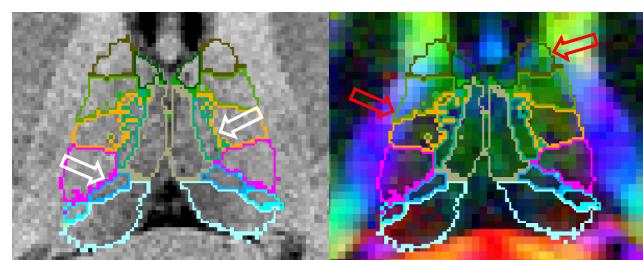


Figure 1: Thalamic segmentation of a T1-weighted structural MRI overlaid on the co-registered T1-weighted image (left) and a co-registered directionally encoded colour FA image (right). High contrast between medial and lateral thalamic regions on structural imaging improves the accuracy of these boundaries (white arrows). However, low contrast between the lateral thalamus and white matter causes over-segmentation into the internal capsule, which can easily be discerned in the colour FA image (red arrows).

1 acquired at 7T (Tourdias et al., 2014; Liu et al., 2020)
2 and on *dMRI* through short-track track density imaging
3 (Basile et al., 2021). In particular, segmentations of 7T
4 white-matter-nulled imaging have been used to generate
5 both multi-atlas segmentation ("THOMAS" Su et al. 2019)
6 and CNN (Umapathy et al., 2021) segmentation algorithms.
7 However, these segmentations do not have the full level
8 of detail present in histological atlases and performance is
9 impacted by changes in acquired contrast, due to domain
10 gap effects for CNNs and poorer registration in multi-atlas
11 segmentation.

12 Aiming to provide detailed segmentations of thalamic
13 nuclei that is robust to changes in MRI acquisition and
14 contrast, we previously constructed a probabilistic atlas of
15 the thalamus and surrounding tissue from manually labelled
16 histology (Iglesias et al., 2018). We then combined this
17 atlas with Bayesian inference methods (Wells et al., 1996;
18 Van Leemput et al., 1999; Ashburner and Friston, 2005; Pohl
19 et al., 2006) to allow segmentation of 25 bilateral histologi-
20 cal labels from sMRI. This approach had the advantage that
21 the intensity model of each label was learned from the target
22 image, allowing the resulting labels to remain stable across
23 acquisition contrasts. However, these segmentations can be
24 less accurate in areas where sMRI shows poor contrast.
25 For example, Fig. 1 shows that our previous method can
26 accurately follow the boundary between groups of medial
27 and lateral nuclei, but the lack of contrast between lateral
28 nuclei and white matter can lead to oversegmentation into
29 the internal capsule.

30 The availability of complementary information from
31 *dMRI* sequences provides a possible avenue for minimising
32 such segmentation errors. An increasing number of
33 large multi-site neuroimaging studies, including the Hu-
34 man Connectome Project (*HCP*) (Van Essen et al., 2013),
35 the Alzheimer's Disease Neuroimaging Initiative (ADNI)
36 (Jack Jr. et al., 2008), and the GENetic Frontotemporal de-
37 mentia Initiative (GENFI) (Rohrer et al., 2015) are acquiring
38 both structural and diffusion MRI. Additionally, use of DTI
39 combined with registration-based segmentation has been
40 proposed for segmentation of the whole thalamus in subjects

1 where T1-weighted MRI contrast is very low (Al-Saady
2 et al., 2022). As can be seen in Fig. 1, dMRI shows good
3 contrast between the thalamus and the adjacent white matter,
4 while structural MRI provides better contrast between the
5 medial nuclei and cerebrospinal fluid (CSF) as well as higher
6 resolution. Therefore, we look towards creating joint models
7 of structural and diffusion MRI, incorporating likelihood
8 models of DTI such as those used in the modelling of white
9 matter fibres (Jian and Vemuri, 2007).

10 We present an extension of our structural Bayesian in-
11 ference segmentation algorithm to incorporate dMRI. We
12 focus on DTI due to the ease of fitting tensors to diffusion-
13 weighted images, even from legacy data or in studies with
14 short acquisitions. We explore our recently proposed diffusion-
15 likelihood model, combining the Dimroth-Scheidegger-
16 Watson (DSW) and Beta distributions (Iglesias et al., 2019).
17 We compare this model to both the Wishart distribution,
18 from fibre modelling literature (Jian and Vemuri, 2007), and the log-Gaussian distribution, influenced by tensor in-
19 terpolation methods (Arsigny et al., 2006). Additionally, we
20 build on our previous histological atlas of the thalamus by
21 adding 45 labels for white matter tracts passing adjacent to
22 the thalamus, allowing the DTI likelihood models to capture
23 the varying directionality of fibers in white matter without
24 becoming sensitive to non-white-matter tissue. The result-
25 ing segmentation method allows constraints to be imposed
26 independently on both the structural and diffusion modelling
27 by including separate shared parameter models, enforcing
28 reflective symmetry, incorporating prior distributions on
29 likelihood parameters, and re-weighting likelihood terms to
30 account for the lower resolution of DTI.

31 This paper is structured as follows. In Section 2 we
32 outline our joint segmentation method. This includes ex-
33 planations of: the general Bayesian inference model; the
34 model fitting and segmentation process; the three likeli-
35 hood models; the atlas and its construction; and general
36 implementation details. In Section 3 we evaluate our joint
37 segmentation method on both high and low resolution data.
38 This evaluation includes: model optimisation and evalua-
39 tion on a population template constructed from both T1-weighted
40 MP-RAGE and DTI images; evaluation of the optimised
41 models on subjects from HCP, providing comparison to
42 manual ground truth and test-retest reliability; and test-
43 retest and indirect evaluation on conventional quality data.
44 Section 4 concludes the paper.

46 2. Bayesian segmentation of brain MRI

47 2.1. Probabilistic model and Bayesian inference

48 Here we outline the theory and implementation of our
49 Bayesian segmentation algorithm. As in existing Bayesian
50 segmentation literature (Van Leemput et al., 1999; Zhang
51 et al., 2001; Ashburner and Friston, 2005; Iglesias et al.,
52 2015; Puonti et al., 2016), our strategy relies on modell-
53 ing the voxel-wise data as observable random variables. These

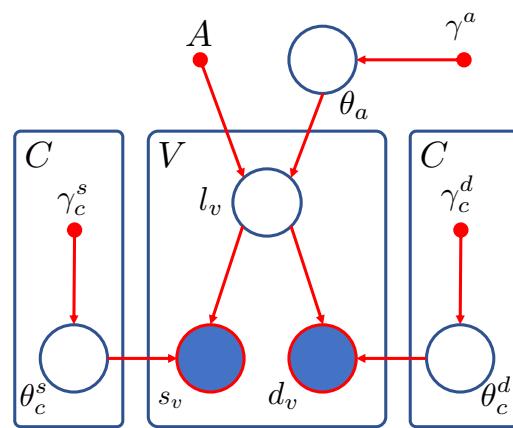


Figure 2: Graphical model of the proposed framework. Larger circles represent random variables with open circles for the hidden variables (θ, l), and shaded circles for the observed variables (s, d). Smaller solid circles are deterministic parameters such as the atlas (A) and encoded prior information (γ). Rectangles indicate replication across voxels (V) or classes (C).

follow a different distribution for each label class in a supplied deformable probabilistic atlas of the volume encompassing the thalamus (Van Leemput, 2009; Iglesias et al., 2018). Both the voxel-data distributions and deformation of the atlas are parameterised by hidden random variables dependent on the subject and image acquisition. Estimating these hidden random variables allows us to generate a voxel-wise probability of membership in each label class (Van Leemput et al., 1999; Ashburner and Friston, 2005). In the Bayesian approach, this is used to construct the posterior probability of a labelling (or segmentation) given paired sMRI and dMRI data.

For the purposes of this method we assume that both the sMRI and dMRI have been registered and resampled to the same grid comprised of voxels indexed by $v \in \{1, \dots, V\}$. We denote the labelling of these voxels by $\mathbf{L} = [l_1, \dots, l_V]$, with $l_v \in \{1, \dots, C\}$ – where C is the number of label classes in our model. Similarly, we construct a matrix $\mathbf{S} = [s_1, \dots, s_V]$ holding vectors of sMRI voxel data, s_v , and matrix $\mathbf{D} = [d_1, \dots, d_V]$ to hold the dMRI voxel data, d_v . We explore different representations of d_v in later sections.

Using this notation and applying Bayes' rule, the posterior probability of a specific labelling for a pair of sMRI and dMRI scans of a subject is:

$$p(\mathbf{L}|\mathbf{S}, \mathbf{D}) \propto p(\mathbf{S}, \mathbf{D}|\mathbf{L})p(\mathbf{L}), \quad (1)$$

and the labelling that maximises Eq. (1) is known as the maximum a posteriori (MAP) estimate for the segmentation. To obtain this MAP estimate we need both the *likelihood* distribution, $p(\mathbf{S}, \mathbf{D}|\mathbf{L})$, of our imaging data given a segmentation, and a *prior* distribution, $p(\mathbf{L})$, generated from prior anatomical knowledge of the thalamus and its surroundings. As these can be used to generate random scans by sampling first from the prior then from the likelihood, segmentation can be thought of as fitting a generative probabilistic forward model to our data and “inverting” it to obtain the labelling.

To make the problem in Eq. (1) tractable, we assume: *i*) that both the likelihood and prior factorise over voxels

Multi-modal thalamic segmentation

1 and *ii*) that the sMRI and dMRI are independent of each
2 other given the labels. The exact graphical model of our
3 framework is shown in Fig. 2. At the top of this model we
4 define the prior distribution on the labels, beginning with a
5 probabilistic atlas A . This atlas is constructed within a reference
6 brain space, meaning it is likely to match the topology
7 of any segmentation subject, but will require deformation to
8 match accurately. The atlas A provides, at each spatial location,
9 the prior probability of observing each neuroanatomical
10 label class. We define A on a deformable tetrahedral
11 mesh, where each vertex has an associated vector of class
12 probabilities, and barycentric interpolation can be used to
13 obtain probabilities at non-vertex locations (Van Leemput,
14 2009). We define a set of parameters, θ^a , that move the
15 mesh nodes to deform the atlas into the space of the target
16 MRI voxel grid, accommodating the anatomical variability
17 across subjects. These parameters are themselves a sample
18 from a distribution that is regularised by setting the stiffness
19 γ^a , preventing folding of the atlas mesh and preserving
20 topology. We then assume that our labelling L is sampled
21 from the categorical distribution over classes defined by the
22 deformed atlas, with each voxel location sampled independently
23 allowing factorisation.

24 Given L we can define the likelihood model for our
25 observed data. We assume that the sMRI and dMRI are
26 conditionally independent from each other and across voxels
27 given the labelling, with s_v and d_v modelled as samples
28 from separate distributions parameterised by θ_c^s and θ_c^d
29 respectively. These hidden parameters are dependent on the
30 corresponding label $l_v = c$. Any prior knowledge on these
31 parameters is encoded in prior distributions controlled by
32 hyperparameters γ_c^s and γ_c^d .

33 Under these assumptions we can define the full joint
34 probability density function (PDF) for Fig. 2 as

$$\begin{aligned} p(S, D, L, \theta | A, \gamma) &= p(S | L, \theta^s) p(D | L, \theta^d) p(L | A, \theta^a) p(\theta | \gamma) \\ &= \left(\prod_{v=1}^V p(s_v | \theta_{l_v}^s) p(d_v | \theta_{l_v}^d) p(l_v | A, \theta^a) \right) \\ &\quad \left(\prod_{c=1}^C p(\theta_c^s | \gamma_c^s) p(\theta_c^d | \gamma_c^d) \right) p(\theta^a | \gamma^a), \end{aligned} \quad (2)$$

35 where $\theta = \{\theta_c^s, \theta_c^d, \theta^a\}$ and $\gamma = \{\gamma_c^s, \gamma_c^d, \gamma^a\}$.

36 With the model described by Fig. 2 and Eq. (2) we can
37 formulate the MAP estimate for our segmentation as

$$\begin{aligned} L_{\text{MAP}} &= \arg \max_L p(S, D | L, A, \gamma) p(L | A, \gamma) \\ &= \arg \max_L \int p(S, D | L, \theta, A) p(L | \theta, A) p(\theta | S, D, A, \gamma) d\theta. \end{aligned} \quad (3)$$

38 However, integrating the joint PDF over the full space of
39 possible parameters θ is intractable. For this reason we make
40 the standard assumption that the posterior distribution of
41 the hidden parameters is heavily peaked around the mode,
42 $p(\theta | S, D) \simeq \delta(\theta - \hat{\theta})$. In this way, we can segment our
43 images by applying Bayes' rule to Eq. (2) and marginalising

1 over the hidden labelling L to obtain these optimal hidden
2 parameters (so called "point estimates"):

$$\hat{\theta} = \arg \max_{\{\theta^a, \theta^s, \theta^d\}} \left[p(\theta^a | \gamma^a) p(\theta^s | \gamma^s) p(\theta^d | \gamma^d) \right. \\ \left. \sum_L p(S, D | L, \theta^s, \theta^d) p(L | \theta^a, A) \right], \quad (4)$$

3 and then optimising L to obtain the MAP estimate

$$L_{\text{MAP}} = \arg \max_L p(S, D | L, \hat{\theta}, A) p(L | \hat{\theta}, A). \quad (5)$$

2.2. Parameter estimation and segmentation

4 The first step in the segmentation process is to estimate
5 the optimal hidden parameters $\hat{\theta}$ from Eq. (4). We begin by
6 formulating the likelihood PDFs for both sMRI and dMRI
7 as mixture models such that each label class is described
8 by mixtures of G structural component distributions and W
9 diffusion component distributions, giving

$$p(s_v | \theta_c^s) = \sum_i g_{c,i} p(s_v | \theta_i^s), \quad p(d_v | \theta_c^d) = \sum_j w_{c,j} p(d_v | \theta_j^d). \quad (6)$$

10 Here, $g_{c,i} \geq 0$ and $w_{c,j} \geq 0$ are mixture weights in the model
11 of label class c indicating the contribution of the i -th sMRI
12 and j -th dMRI component's distribution to the appearance of
13 the class in the respective modality. These distributions are
14 parameterised by θ_i^s and θ_j^d , respectively, with $i \in 1, \dots, G$
15 and $j \in 1, \dots, W$. In both cases the sum over the component
16 weights for a given class must be equal to one, $\sum_i g_{c,i} = 1$
17 and $\sum_j w_{c,j} = 1$. This mixture model formulation provides
18 a high degree of flexibility, allowing us to specify *a priori*
19 which label classes may be modelled jointly by constraining
20 specific weights to 0 or 1 while others are allowed to vary.

21 Combining Eq. (6) with Eqs. (2) and (4) and taking
22 logarithms we can then obtain an objective function to be
23 optimised with respect to the distribution parameters,

$$\begin{aligned} O(\theta | S, D, A, \gamma) &= \log p(\theta^a | \gamma^a) + \sum_i^G \log p(\theta_i^s | \gamma_i^s) + \sum_j^W \log p(\theta_j^d | \gamma_j^d) \\ &\quad + \sum_v^V \log \sum_c^C p(l_v^c | A, \theta^a) \left[\sum_i^G g_{c,i} p(s_v | \theta_i^s) \right] \left[\sum_j^W w_{c,j} p(d_v | \theta_j^d) \right]. \end{aligned} \quad (7)$$

24 To optimise Eq. (7) we adapt the approach proposed by
25 Puonti et al. (2016). In this approach the atlas deformation
26 parameters and likelihood parameters are optimised
27 iteratively in a coordinate ascent scheme, with each being
28 optimised while the other is fixed. The optimisation of the
29 θ^a is performed using a standard conjugate gradient operator
30 with the deformation prior $p(\theta^a | \gamma^a)$ taking the form of
31 the penalty term defined by Ashburner et al. (2000). The
32 likelihood parameters θ^s and θ^d are then optimised using
33 a Generalised Expectation Maximisation (GEM) algorithm
34 (Dempster et al., 1977; Van Leemput et al., 1999), iterating
35 between expectation (*E*) and Maximisation (*M*) steps.

1 **E step:** In the E step, we build a lower bound $Q(\theta)$ for the
 2 objective function in Eq. (7) using Jensen's inequality:

$$Q(\theta) = \log p(\theta^a | \gamma^a) + \sum_i^G \log p(\theta_i^s | \gamma_i^s) + \sum_j^W \log p(\theta_j^d | \gamma_j^d) \\ + \sum_{v,c,i,j} q_v^{c,i,j} \log \left[p(l_v^c | A, \theta^a) p(s_v | \theta_i^s) p(\mathbf{d}_v | \theta_j^d) \right] \\ - \sum_{v,c,i,j} q_v^{c,i,j} [\log q_v^{c,i,j} - \log g_{c,i} - \log w_{c,j}]. \quad (8)$$

3 Here l_v^c indicates the event that the voxel label $l_v = c$ and
 4 $q_v^{c,i,j}$ is a soft segmentation at the current parameter estimates
 5 indicating the combination of class c , sMRI distribution i
 6 and dMRI distribution j :

$$q_v^{c,i,j} = \frac{g_{c,i} w_{c,j} p(l_v^c | A, \theta^a) p(s_v | \theta_i^s) p(\mathbf{d}_v | \theta_j^d)}{\sum_{\{v,i,j\}} g_{c,i} w_{c,j} p(l_v^c | A, \theta^a) p(s_v | \theta_i^s) p(\mathbf{d}_v | \theta_j^d)}. \quad (9)$$

7 **M step:** In the generalised M step we attempt to increase
 8 the bound $Q(\theta)$ in Eq. (8). We note that the two sets of distri-
 9 bution parameters θ_i^s and θ_j^d can be optimised individually,
 10 as they make independent contributions to the bound:

$$Q_s(\theta_i^s) = \log p(\theta_i^s | \gamma_i^s) + \sum_v^V \left[\sum_{c,j} q_v^{c,i,j} \right] \log p(s_v | \theta_i^s), \quad (10)$$

$$Q_d(\theta_j^d) = \log p(\theta_j^d | \gamma_j^d) + \sum_v^V \left[\sum_{c,i} q_v^{c,i,j} \right] \log p(\mathbf{d}_v | \theta_j^d). \quad (11)$$

11 These contributions can then be optimised using either
 12 closed form solutions or numerical methods, depending on
 13 the distribution used as we will describe in Section 2.3.
 14 Finally we can calculate the new optimal weightings as

$$g_{c,i} = \frac{\sum_{\{v,j\}} q_v^{c,i,j}}{\sum_{\{v,i,j\}} q_v^{c,i,j}} \quad w_{c,j} = \frac{\sum_{\{v,i\}} q_v^{c,i,j}}{\sum_{\{v,i,j\}} q_v^{c,i,j}} \quad (12)$$

15 **Segmentation:** The mesh deformation and likelihood
 16 parameter optimisation steps are repeated alternately until the
 17 objective function in Eq. (7) has converged. At this point,
 18 we note that the formulation of the posterior factorises over
 19 voxels and the posterior probability of each class may be
 20 found by summing over the soft segmentations $q_v^{c,i,j}$. Hence
 21 the final MAP estimate segmentation is given by

$$\hat{l}_v = \arg \max_c \sum_{i=1}^G \sum_{j=1}^W q_v^{c,i,j}. \quad (13)$$

22 **2.3. Likelihoods**

23 So far, we have outlined the Bayesian framework and
 24 segmentation process without specifying the likelihood
 25 models used for both sets of MRI data. The steps outlined
 26 above are not affected by the choice of distributions used.
 27 Here we provide an overview of the distributions used to
 28 model the sMRI and dMRI data, including the likelihood
 29 term and, where applicable, the prior over its parameters.
 30 Detailed equations for the calculation of PDF values as well
 31 as the optimisation of model parameters, θ , may be found in
 32 Section S.1 of the supplement.

2.3.1. Structural MRI model

To model the sMRI intensities, we follow the Bayesian brain MR segmentation literature and use a mixture of Gaussian intensity distributions (Ashburner and Friston, 2005; Zhang et al., 2001; Van Leemput et al., 1999). In this model the intensity values for each structural modality are held in the vector s_v and the model parameters θ_i^s are the mean and covariance, $\{\mu_i, \Sigma_i\}$, of the structural mixture component i . We choose to use the natural conjugate prior, the Normal-Inverse-Wishart distribution, on these Gaussian parameters. The likelihood and prior distributions are therefore

$$p(s_v | \theta_i^s) \sim \mathcal{N}(\mu_i, \Sigma_i), \quad p(\mu_i, \Sigma_i | \gamma_i^s) \sim \mathcal{NIW}(\mathbf{M}_i^s, n_i^s, \Psi_i^s, v_i^s), \quad (14)$$

where $\mathbf{M}_i^s, n_i^s, \Psi_i^s$ and v_i^s encode any prior knowledge we may have on the structural distribution. Formulations for the structural PDFs and closed form solutions to the parameter M step parameter optimisations can be found in Section S.1.1 of the supplement.

2.3.2. Diffusion MRI models

To model the dMRI data, we consider distributions over tensors estimated with DTI. Even though higher-order models can be used with modern dMRI acquisitions, using DTI models ensures that our method is compatible with virtually every dMRI dataset, including huge amounts of legacy data. In this work, we compare two competing models, based on the Wishart and Gaussian distributions, to our previously-proposed DSW-beta distribution (Iglesias et al., 2019).

Wishart: Following existing white matter fibre modelling literature, we look to the Wishart distribution (Jian and Vemuri, 2007). DTI produces at each voxel a covariance matrix describing the displacements of water molecules in the voxel. Therefore, the natural conjugate prior for these tensors is an Inverse-Wishart distribution. We use this in combination with a Gamma distribution on the degrees of freedom parameter (Görür and Rasmussen, 2010), with the effect of lowering the degrees of freedom and increasing the breadth of the resulting Wishart distributions. In this model, we define \mathbf{d}_v as the inverse of the diffusion tensor T_v . We then use the Wishart and Gamma distributions to model \mathbf{d}_v and θ_j^d :

$$\mathbf{d}_v \sim \mathcal{W}(n_j^d, V_j^d), \quad (n_j^d - 2)/2 \sim \Gamma(\alpha, \beta), \quad (15)$$

where α and β are set to 0.5 and 1.5 respectively to provide a non-informative prior. Formulations for the Wishart PDFs and the optimisation problem in the M step can be found in Section S.1.2 of the supplement.

Log-Gaussian: This model is motivated by literature on the interpolation of DTI volumes. Direct interpolation of DTI can lead to swelling of the ellipsoids representing the diffusion tensors, but interpolating in the log domain reduces this effect (Arsigny et al., 2006; Dryden et al., 2009). For this reason, and noting that the DTI tensors, T_v , are symmetric with only six independent variables, we define \mathbf{d}_v as a vector

$$\mathbf{d}_v = P \text{vec}(\log T_v), \quad \text{vec}(\log T_v) = P^\top \mathbf{d}_v, \quad (16)$$

Multi-modal thalamic segmentation

1 where P is a constant 6×9 matrix (values listed in supplement)
2 designed with the constraint that

$$\|\log(T_1) - \log(T_2)\|_F^2 = \|\mathbf{d}_1 - \mathbf{d}_2\|_2^2, \quad (17)$$

3 and therefore interpolation of the vectors \mathbf{d}_v is equivalent
4 to interpolation of the tensors in the log domain. In this
5 formulation the natural distribution to choose based on the
6 distance metric in Eq. (17) is a Gaussian distribution with a
7 scalar variance

$$\mathbf{d}_v \sim \mathcal{N}(\mathbf{m}_j^d, \sigma_j^d). \quad (18)$$

8 We then define uniform priors on both \mathbf{m}_j^d and σ_j^d due to
9 the difficulty in informing these parameters a priori. Formulations
10 for the log-Gaussian PDFs and the optimisation problem in the M step can be found in Section S.1.3 of the
11 supplement.

12 **DSW-beta:** This model is a custom distribution proposed
13 in our prior work (Iglesias et al., 2019) motivated by a desire
14 to lower the dimensionality of \mathbf{d}_v , leading to a reduction
15 in extreme values of the likelihood that may overwhelm
16 the prior. Here only the fractional anisotropy (FA), f_v , and
17 the principal eigenvector, ϕ_v , of the tensor T_v are modelled
18 so that $\mathbf{d}_v = \{f_v, \phi_v\}$. In this approach, we use the two
19 parameter Beta distribution to model the FA as it is able to
20 model both the location and dispersion of signals in the range
21 $[0, 1]$. We then use the DSW distribution to model ϕ_v .

22 This DSW distribution has the advantages of symmetry
23 and simplicity. As DSW is antipodally symmetric, it
24 accommodates the directional invariance of dMRI (Zhang
25 et al., 2012). It is also rotationally symmetric around a mean
26 direction and its opposite $\{\psi, -\psi : \|\psi\| = 1\}$, with a
27 dispersion around the mean parameterized by a concentration
28 κ (Mardia et al., 2000). This κ allows us to incorporate
29 the higher directional dispersion in voxels with lower FA
30 by multiplying the component specific concentration by the
31 voxel FA to give an effective concentration for each voxel.
32 The likelihood distribution in this formulation of the dMRI
33 is therefore a joint DSW-beta distribution

$$f_v \sim \mathcal{B}(\alpha_j^d, \beta_j^d), \quad \phi_v \sim \mathcal{DSW}(\psi_j^d | f_v \kappa_j^d). \quad (19)$$

35 Formulations for the DSW-beta PDFs and M step can be
36 found in Section S.1.4 of the supplement.

37 2.4. Prior distribution: an improved probabilistic 38 atlas of the thalamus

39 In Iglesias et al. (2018), we presented a highly detailed
40 probabilistic atlas of the human thalamus built from a combination
41 of *in vivo* MRI and histology. The spatial distribution
42 of the thalamic nuclei was learnt from manual delineations
43 drawn on 3D reconstructed histological sections from 12
44 specimens (Fig. 3a), whereas 39 MRI scans with manual
45 delineations (Fischl et al., 2002) were used to learn the
46 distribution of surrounding tissue (Fig. 3b). Direct use of
47 this atlas (Fig. 3d) in our new framework is not ideal, as the
48 cerebral white matter was modelled using only two classes –
49 one per hemisphere. While such a parsimonious model with

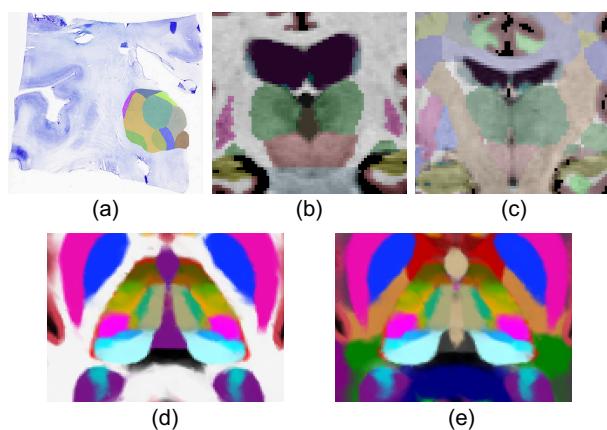


Figure 3: (a-c) Types of segmentations used to build the atlas. (a) Coronal histological section of the thalamus, with manual delineations of the nuclei. (b) Coronal slice of an *in vivo* T1-weighted MRI scan, with manual delineations for whole brain structures. (c) Similar coronal slice of one of the new 16 cases, with the white matter subdivided into tracts. (d-e) Corresponding axial slices of the previous and updated probabilistic atlases. The original atlas (d) was trained with segmentations like the ones in (a-b), while the new atlas used (a-c).

1 a single component is adequate for modelling the unimodal
2 distribution of white matter intensities in sMRI, it is largely
3 insufficient to model the dMRI orientations. The distribution
4 over white matter voxels is highly multimodal due to the
5 variety of fibre tracts that traverse this tissue in different
6 orientations.

7 In principle we could model such a complex distribution
8 using a mixture model with many components. However,
9 such an approach is likely to fail, as some of these com-
10 ponents may end up modelling non-white-matter tissue.
11 Instead, we have refined our atlas by subdividing the white
12 matter surrounding the thalamus into 45 tracts. To achieve
13 this, we complemented the training data in Iglesias et al.
14 (2018) (12 *ex vivo* thalami and 39 *in vivo* whole brains)
15 with *in vivo* sMRI/dMRI data from 16 additional subjects,
16 that were labelled manually as part of an update (Maffei
17 et al., 2021) to the TRACULA (TRacts Constrained by
18 UnderLying Anatomy) package distributed with FreeSurfer
19 (Yendiki et al., 2011).

20 The TRACULA training set (16 healthy adults from
21 the publicly available MGH-USC HCP; Fan et al. 2016)
22 consisted of dMRI data, acquired using 512 directions at a
23 maximum b-value of 10,000 s/mm² with 1.5 mm isotropic
24 spatial resolution, and sMRI T1-weighted data, acquired
25 with an MPRAGE sequence at 1 mm isotropic resolution.
26 Cortical parcellations and subcortical segmentations, in-
27 cluding the whole thalami and cerebral white matter (left
28 and right), were obtained through FreeSurfer (Dale et al.,
29 1999; Fischl et al., 1999, 2002, 2004). Whole-brain prob-
30abilistic tractograms were generated for each subject using
31 constrained spherical deconvolution approaches (Tax et al.,
32 2014; Jeurissen et al., 2014) and streamlines used to man-
33 ually label 42 white matter tracts through a combination

Multi-modal thalamic segmentation

1 of inclusion and exclusion criteria (Maffei et al., 2021).
2 Resulting tractograms were transformed to the sMRI of the
3 subject using a boundary-based, affine registration method
4 (Greve and Fischl, 2009) and converted into visitation maps.
5 These soft segmentations were spatially smoothed with a
6 Gaussian kernel ($\sigma = 2\text{mm}$). For each white matter voxel
7 in the FreeSurfer subcortical segmentation, we replaced its
8 label by the tract with the highest probability (unless such
9 probability was below 5%), dividing the white matter into
10 42 tracts and a generic white matter class (Fig. 3c).

11 The three types of segmentations (Fig. 3a-c) were used
12 to rebuild the atlas, using a technique that enables combining
13 labellings with different levels of detail (Iglesias et al., 2015).
14 The new atlas (Fig. 3e) is almost identical to the old one
15 (Fig. 3d), but now includes more specific subclasses in
16 the white matter. As a last adjustment, we subdivided the
17 anterior commissure and other tracts comprising the corpus
18 callosum, while excluding tracts not passing adjacent to
19 the thalamus. This resulted in 45 final labels for the white
20 matter tracts. Each of these subclasses can be modelled
21 either with unimodal distributions or mixtures with very few
22 components, effectively preventing the modelling of non-
23 white-matter tissue.

2.5. Implementation details

2.5.1. Data preparation

26 We assume that the sMRI has been processed with
27 FreeSurfer, which yields a bias field corrected image and a
28 whole brain segmentation (*aseg.mgz*, Fischl et al. 2002). The
29 labels in *aseg.mgz* are used to initialise both the atlas deformation
30 (Iglesias et al., 2015, 2018) and hyperparameters in
31 the structural prior in Eq. (14). In practice the hypermean
32 \mathbf{M}_i^s is estimated from the median of the relevant label in this
33 initial coarse segmentation, and n_i^s relates to the number of
34 voxels used in estimating \mathbf{M}_i^s . However, it is more difficult
35 to robustly inform prior distributions of the covariance, so
36 we set both Ψ_i^s and ν_i^s to zero to provide a non-informative
37 prior, giving the set of prior parameters $\gamma_i^s = \{\mathbf{M}_i^s, n_i^s\}$.

38 We also assume that the source dMRI has been put
39 through the preprocessing stages of TRACULA (Yendiki
40 et al., 2011; Maffei et al., 2021). This includes FSL's
41 eddy current and subject motion correction (Andersson
42 and Sotiroopoulos, 2016) before fitting the tensor model.
43 Additionally, we identify DTI voxels with poor fits as those
44 with tensors that have negative eigenvalues or FA outside
45 the range [0, 1]. These are replaced by a local average tensor
46 constructed by convolution of the log space tensors with a
47 3D Gaussian kernel. These cleaned tensors are converted to
48 the log domain (Arsigny et al., 2006) before interpolation to
49 the voxel grid of the sMRI.

2.5.2. Mixture model specification

50 The assignment of component distributions to label
51 classes is one of the modelling choices that must be made
52 before segmentation. We assign structural and diffusion
53 components independently for each label class, defining
54 what we will call the structural mixture model (*sMM*) and
55 diffusion mixture model (*dMM*) respectively. In practice,
56

1 this constrains most weights $g_{c,i}$ and $w_{c,j}$ to 0 or 1, with a
2 single component distribution often shared between groups
3 of labels. However, we do allow for many-to-many relations
4 between the label-classes and components. For example,
5 allowing the structural appearance of the CSF label
6 to be modelled by two Gaussian components, one for "clean"
7 CSF that is also used to model ventricle labels and one for
8 "messy" CSF that is shared with the choroid plexus.

9 For class likelihoods composed of multiple distributions,
10 the non-zero weights are set to be equal for the first E step and
11 initial component parameters are obtained by use of k-means
12 clustering. Details of this clustering for each likelihood
13 formulation can be found in Section S.3 of the supplement,
14 while optimisation of the default sMM and dMM definitions
15 is performed in Section 3.2.

2.5.3. Reflective symmetry

16 Another regularising constraint we impose on our seg-
17 mentation is reflective symmetry of contralateral structures.
18 Classes in one hemisphere share structural distributions
19 with the corresponding classes in the opposite hemisphere.
20 However, in dMRI we assume the average ellipsoids de-
21 scribed by tensors from two contralateral structures should
22 be reflections of each other in the median plane. As the
23 head is never positioned in a perfect alignment with the
24 scanner coordinate system, we optimise for a vector normal
25 to the plane of reflection, \mathbf{r} , initially assumed to be parallel
26 to the left-right axis of the voxel grid. Prior to the M step,
27 we substitute reflected distribution parameters to the bound
28 in Eq. (8) and formulate the contribution of \mathbf{r} , producing
29 an objective function that is fourth order in \mathbf{r} with known
30 first and second derivatives. This objective can be optimised
31 using an interior-point method under the constraint that
32 $\|\mathbf{r}\| = 1$. We then jointly fit parameters for corresponding
33 component distributions in the left and right hemispheres.
34 Detailed formulations for the reflection optimisation and
35 joint distribution fitting can be found in Section S.1 of the
36 supplement.

2.5.4. Likelihood adjustment

37 Our model assumes that the resolutions of the dMRI
38 and sMRI are identical. While datasets such as the HCP
39 deviate from this assumption to a lesser degree, conventional
40 quality datasets have much lower resolution for the dMRI
41 in particular, for example T1-weighted images are typically
42 acquired with each voxel dimension at approximately 1 mm
43 while dMRI voxel dimensions can approach 2.5 mm in
44 each direction. As we resample to the resolution of the
45 sMRI, more dMRI voxels are used in likelihood parameter
46 estimation than are available from the source imaging, which
47 leads to overfitting of the dMRI. In practice, we counteract
48 this effect by downplaying the weight of the dMRI voxels in
49 the objective function by a factor ϵ equal to the ratio of voxel
50 sizes between dMRI and sMRI. Further details can be found
51 in Section S.2 of the supplement.

3. Experiments and Results

To quantitatively evaluate our method and compare between the three likelihood formulations we performed experiments using co-registered sMRI and dMRI from three datasets. In Section 3.1 we generate a population template from HCP subjects, and use it to identify manually segmentable labels corresponding to groups of labels from our histological atlas. In Section 3.2 we use this template to tune our method in a process of model selection. In Section 3.3 we evaluate application of our method to high resolution dMRI on subjects from HCP, including comparisons to manual segmentations and test-retest reliability. Finally, in Section 3.4 we evaluate application of our method to conventional quality dMRI. This includes test-retest reliability on images acquired locally at the University College London Dementia Research Centre (*UCL DRC*) and indirect evaluation on subjects with underlying pathologies by testing our method's ability to distinguish between healthy controls and subjects with AD from the ADNI dataset.

In the following experiments, when comparing regions of interest (ROIs) corresponding to the same label in two separate segmentations we use the Dice Similarity Coefficient (*DSC*) and 95th percentile of Hausdorff distance (*95HD*). For two ROIs X and Y these are defined as

$$DSC(X, Y) = \frac{2\|X \cap Y\|}{\|X\| + \|Y\|}, \quad (20)$$

$$95HD(X, Y) = \max(d_{95}(X, Y), d_{95}(Y, X)), \quad (21)$$

where $\|\cdot\|$ indicates the volume of the ROI and $d_{95}(X, Y)$ is the 95th percentile of the set of distances between points on the ROI boundaries, $\{d_x = \min_{y \in S_Y} |x - y|\}_{x \in S_X}$

3.1. Population template and manual labels

When evaluating segmentation methods for medical images, it is common practice to compare the resulting label maps to a gold standard, usually obtained from manual delineation by a trained rater. However, manual delineation of 52 histological labels on *in vivo* MRI is infeasible, as many of the boundaries between are invisible at $\sim 1\text{mm}$ resolution. Manual segmentation protocols for larger groups of thalamic regions (with fewer labels) exist in the literature (Tourdias et al., 2014), but their anatomical definitions are incompatible with those of our histological labels, introducing bias and preventing direct and fair comparison. In this study, our goal is to compare the performance of our tool with a gold standard that is based on our 52 histological labels and informed by both sMRI and dMRI contrast. For this reason, we adapted these labels to define our own manual segmentation criteria for thalamic labels that can be accurately visualised and segmented on a combination of T1-weighted MPRAGE and directionally-encoded colour FA (*DEC-FA*); when labels of smaller thalamic nuclei were not identifiable from the intensity and contrast of the MRIs, these labels were combined and grouped together, so that the boundaries of the original 52 histological atlas labels can be easily matched and compared.

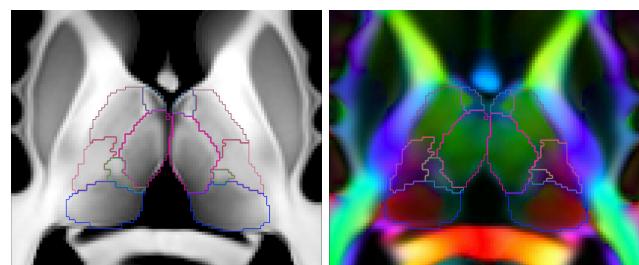


Figure 4: Axial views of the T1-weighted (left) and DEC-FA (right) population templates of the thalamus, overlaid with the outlined labels obtained by manual segmentation. Manually segmented label colour maps are given in Table 1.

Table 1

Summary of the label merging operations used to generate the manually segmented labels from histological atlas nuclei, and groupings of manual labels used for evaluation. Displayed colours follow the convention used in figures throughout this manuscript. Abbreviation definitions are listed in Section S.4 of the supplement.

Grouping	Manual label	Histological atlas labels										
Ant-Lat	Anterior	AV										
	Dorsal	LD	LP									
	Lat-Rostral	VA	VAmc	VLa	Vlp	VM						
Lat-Caudal	Lat-Caudal	VPL										
	Intralaminar	CeM	Pc	Pf	MV(Re)	Pt						
Medial	Int-Imnr-Post	CM										
	Medial	CL	MDI	Mdm								
	Pulvinar	L-Sg	PuA	Pul	PuL	PuM						
Posterior	LGN	LGN										
	MGN	MGN										

The first step in defining these criteria was to create a high resolution template using 500 subjects from the WashU-UMN HCP dataset (Van Essen et al., 2013) and an unbiased template construction method (Joshi et al., 2004). We used three channels in the registration: T1-weighted intensity, T2-weighted intensity, and FA. In order to include directional information in the template, we used the final set of registrations to align and average the DTI tensors in the log domain. The resolution of the template is equal to the resolution of the HCP sMRI data, i.e., 0.7mm isotropic. Slices from the template are shown in Fig. 4.

As a second step to define the gold standard for comparison, we registered the histological atlas to the template, producing a preliminary segmentation of 52 separate thalamic labels. This preliminary segmentation was then manually refined by an anatomy expert (JA, assisted by MB), to correct any anatomical errors from registration, and to combine those thalamic regions which were not reliably identifiable from the multi-modal template into labels which represent larger thalamic groups. This process resulted in a set of 10 bilateral labels that were manually identifiable from the template. The template labels are used in Section 3.2 to aid in tuning our method, and criteria for these labels are used in Section 3.3 to generate gold standard segmentations for comparison. However, the reduced contrast and resolution of scans for individual subjects resulted in increased ambiguity

Multi-modal thalamic segmentation

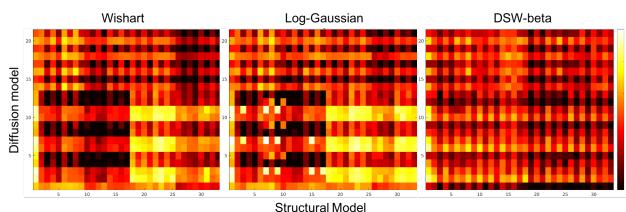


Figure 5: TOPSIS fitness plots for combinations of structural (horizontal axis) and diffusion (vertical axis) grouping models. Plots are displayed for Wishart, Log-Gaussian and DSW-beta likelihood models. A mapping from model numbers to parameter groupings is provided as a spreadsheet in the supplementary material.

for some boundaries, hence we further combine our 10 labels into a final set of 5 coarser groupings that are visually identifiable *in vivo*, enabling evaluation without biasing results. Manual labels for the template can be seen in Fig. 4 and the correspondences between the evaluation groupings, manually segmented labels and original histological atlas labels can be seen in Table 1.

3.2. Model selection

Practical implementation of the proposed framework requires decisions on how to share the sMM and dMM parameters (Section 2.5.2), which amounts to a model selection problem. In principle, our generative models enables the computation of the so-called model evidence, which enables comparison of models with different number of parameters. While theoretically appealing, computing this evidence requires marginalisation over all parameters, which leads to intractable integrals that require approximations. Instead, we selected the sMM/dMM groupings with a combination of prior knowledge and a systematic approach called “Technique for Order Preference by Similarity to Ideal Solution” (TOPSIS) Behzadian et al. (2012).

Structural groupings. In our previous work, we used two Gaussian components to model the contrast difference between medial and lateral classes (Iglesias et al., 2018). Here, we added a third Gaussian modelling the medial portion of the medial pulvinar (PuM) nucleus, which has a structural appearance closer to grey matter compared with the lateral portion of the PuM. We then compared the atlas prior and histograms of the template volumes to identify 33 possible sMMs grouping nuclei into three component distributions, which were considered by TOPSIS (detailed below).

Diffusion groupings. In Section 3.1 we defined 10 labels for each thalamus that are manually identifiable from combined sMRI and DEC-FA. However, inspection of the dMRI tensors within these regions found greater heterogeneity in some regions than in others. As additional borders within these labels could not be confidently matched with boundaries in the histological atlas, we examined multiple options for combining histological nuclei into larger structures to be fit with a component distribution. Including these additional boundaries, and allowing for the possibility of bimodal

histograms for some labels, we arrived at 21 possible dMMs, grouping nuclei into between 11 and 13 component distributions.

TOPSIS. To optimise the choice of sMMs and dMMs in a systematic fashion, we tested each possible combination of sMM and dMM parameter groupings on the population template. We then evaluated these models by comparing Dice scores and 95HD for the whole thalamus as well as the “grouping” and “manual label” regions listed in Table 1. We then used TOPSIS to create a single, normalised fitness score for each combination of shared parameter specifications allowing them to be ranked. The resulting scores for each likelihood model are shown in Fig. 5. The chosen models are provided in a spreadsheet in the supplementary material as well as descriptions of all candidate models.

3.3. Application to high resolution dMRI

Having individually tuned the mixture models and defined a manual protocol corresponding to our histological labels, the obvious next step is to assess the performance of our joint segmentation method on HCP quality data. A comparison of our joint segmentation to both the FreeSurfer whole thalamus segmentation (*aseg.mgz*) and our previous structural-only method are shown in Fig. 6. This figure shows each segmentation overlaid on both the T1-weighted sMRI and the DEC-FA for two healthy subjects².

In both subjects the whole thalamus *aseg* segmentation, used as an initialisation for both Bayesian methods, shows obvious errors when overlaid on the DEC-FA, with more extreme over-segmentation for subject 2. In subject 1 the structural-only segmentation appears to compensate for these errors and provides an improved exterior boundary. However, our joint method shows marked improvement in the agreement of internal boundaries with colours displayed in the dMRI (solid arrows) as well as a smaller improvement in the exterior boundary. This effect is much more pronounced in subject 2, where the initial over-segmentation of the thalamus propagates to the structural-only method but is corrected by the joint method (arrow outlines).

Such observations provide compelling qualitative evidence for the efficacy of our new method. However, to fully evaluate its usefulness we must quantitatively assess both accuracy and repeatability.

3.3.1. Direct evaluation with manual ground truth

To provide a quantitative measure of segmentation quality, our anatomy expert (JA, assisted by MB) manually segmented images for 10 randomly selected subjects from the WashU-UMN HCP dataset (Van Essen et al., 2013) using the protocol outlined in Section 3.1. The manual segmentations were performed using a combination of T1-weighted and DEC-FA at a 1.25 mm isotropic resolution, corresponding

²The joint segmentation shown here uses our DSW-beta likelihood model and the structural method has been optimised for the HCP dataset by tuning of the stiffness parameter. For a visual comparison of all likelihood models and the default structural segmentation please see the Section S.5 of the supplement.

Multi-modal thalamic segmentation

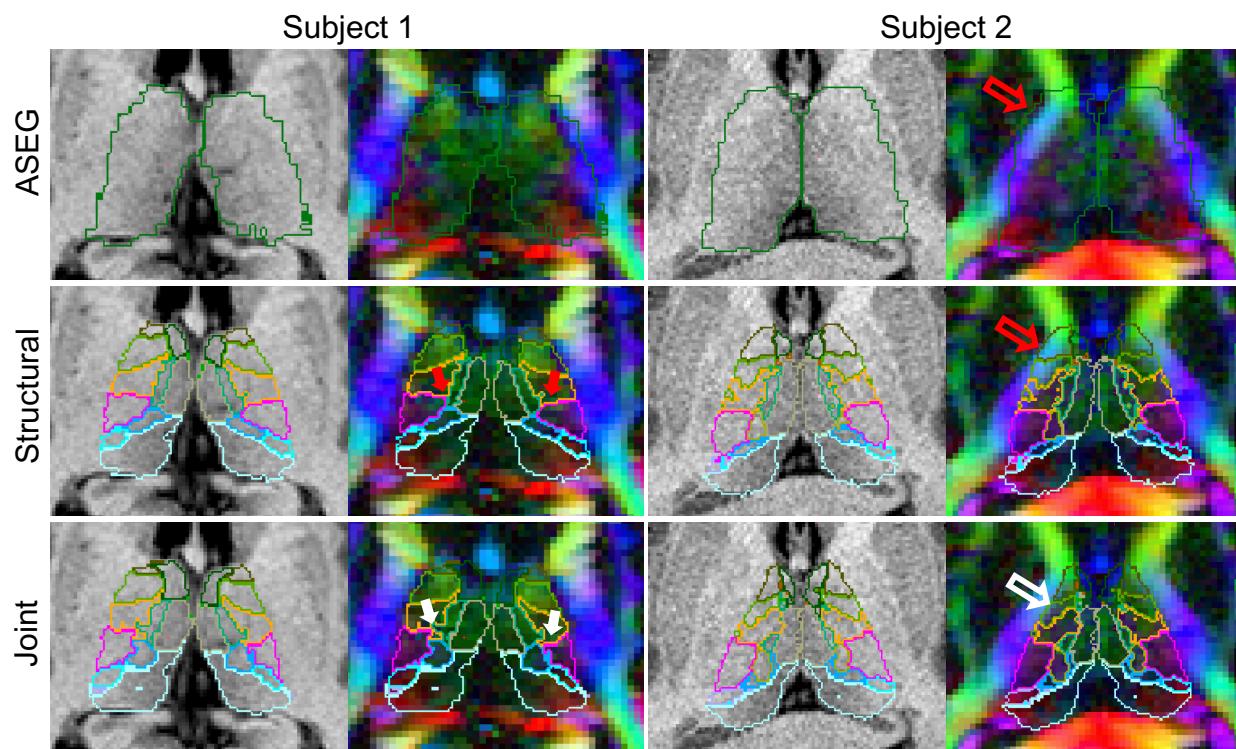


Figure 6: Comparison of thalamic segmentations generated from FreeSurfer's recon-all (aseg.mgz), structural and joint (DSW-beta) Bayesian segmentation on two HCP subjects. Arrows have been overlaid to indicate significant instances of correct (white) and incorrect (red) identification of both internal (solid) and external (outline) boundaries.

1 to the native resolution of the diffusion data in HCP. We
2 generated segmentations for these subjects using each of
3 the three joint likelihood implementations from Section 2.3
4 as well as our previously published structural-only imple-
5 mentation³ (Iglesias et al., 2018). These automated seg-
6 mentations, which have the resolution of the structural scans
7 (0.7 mm), were resampled to 1.25 mm isotropic resolution
8 and compared with the ground truth using DSC and 95HD.
9 Dice scores and 95HD for the five groupings (in column one
10 of Table 1) and the whole thalamus are shown in Fig. 7.

11 In general, both the DSC and 95HD plots follow similar
12 trends. The median Dice scores for the whole thalamus in
13 structural-only, Wishart and Log-Gaussian implementations
14 were 0.88 with a small increase to 0.89 for DSW-beta. Simi-
15 larly the 95HD for all methods was between 2.3 and 2.5 mm
16 or equivalent to approximately 2 voxels on the manual seg-
17 mentations. This contrasts to the marked improvement in the
18 exterior boundary for subject 2 in Fig. 6. As subject 2 was not
19 selected for manual segmentation, the direct comparisons in
20 Fig. 7 indicate that joint segmentation may only provide a
21 small improvement in exterior boundary where the structural
22 segmentation has worked well, but that it is more robust to
23 errors in initialisation and atypical cases.

24 Of more interest are the interior boundaries. In nearly
25 all labels the joint methods outperform structural-only with

1 lateral-caudal class showing an improvement of 10 Dice
2 points. This can be seen in Fig. 6 where the solid arrows
3 indicate this changes in this boundary for subject 1. The
4 only label class where the structural method outperforms our
5 joint implementation is the medial class. This is expected
6 as a medial-lateral contrast change is modelled explicitly in
7 the structural-only method. However, the difference is small
8 with a median DSC of 0.72 in structural compared to 0.67
9 for the joint methods and comparable 95HD measurements
10 in this class.

11 There is comparatively little difference between the three
12 diffusion likelihood implementations. The Wishart and Log-
13 Gaussian implementations show the most similar results,
14 while in the DSW-beta implementation small decreases in
15 accuracy of the intralaminar and posterior classes are offset
16 by improvements in the antero-lateral classes and whole
17 thalamus exterior.

3.3.2. Test-retest reliability analysis

18 In order to assess the test-retest reliability of the method
19 (a crucial feature in large scale, multi-centre studies), we
20 segmented images from 110 HCP subjects using two differ-
21 ent sets of DTI images for each subject – one based on the
22 $b=1000 \text{ s/mm}^2$ shell and one based on the $b=2000 \text{ s/mm}^2$
23 shell – and compared the outputs. While the results of such
24 an experiment are optimistic when compared to experiments
25 in which images are acquired with multiple scanners, it does
26 enable thorough comparison within the same dataset; test-
27 retest experiments with multiple acquisitions are described
28 in Section 3.4.1 below.

³It should be noted that, to ensure a fair comparison of joint and single channel approaches, the mesh stiffness parameter of the structural implementation was modified to match the joint model that had been developed on the HCP dataset. This improved both the DSC and 95HD structural results compared to the default FreeSurfer distribution.

Multi-modal thalamic segmentation

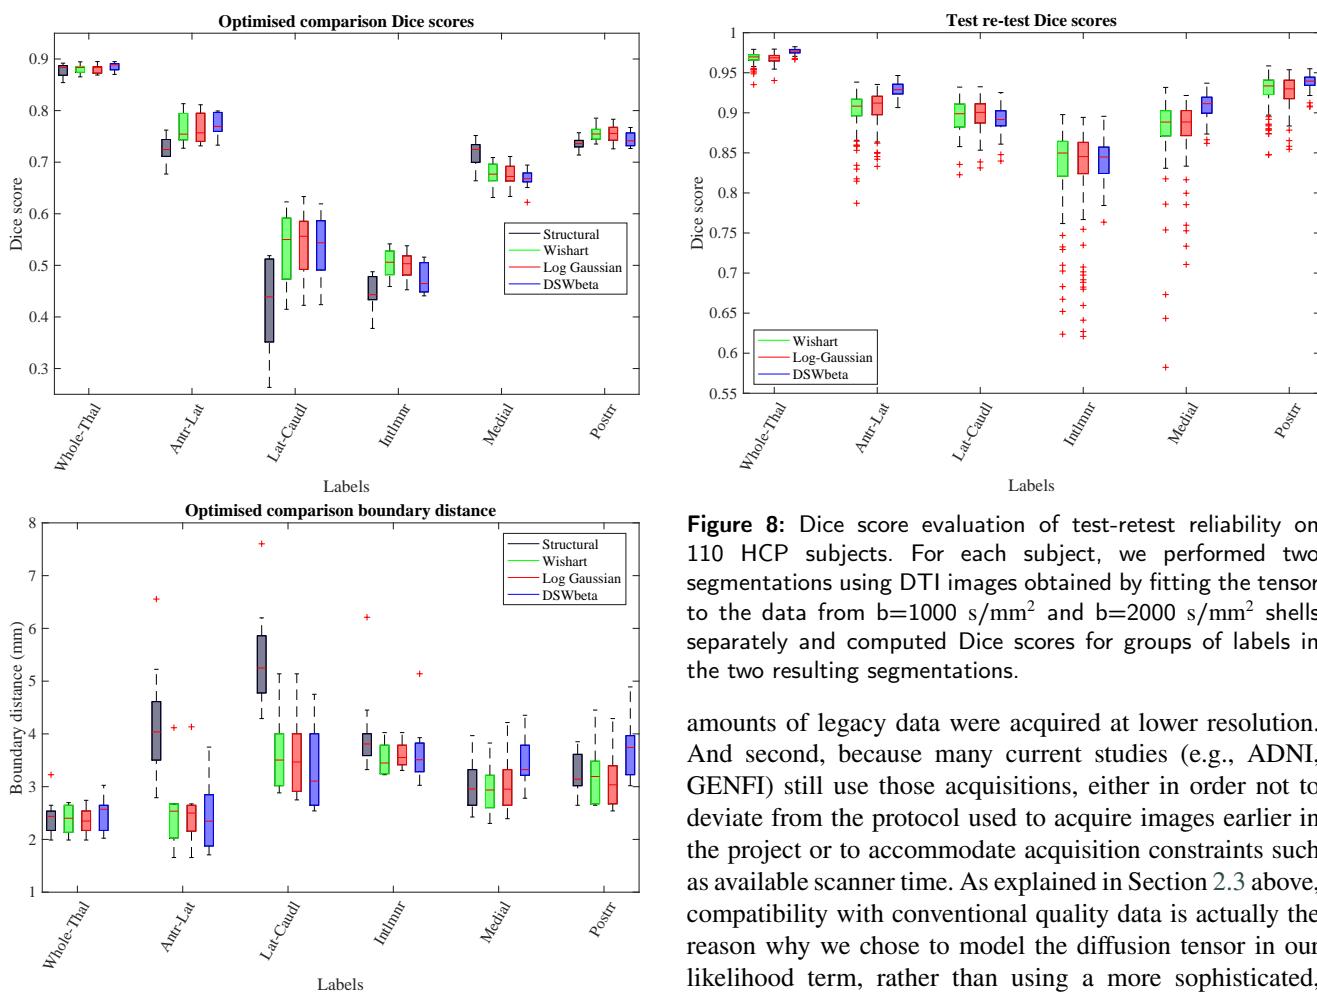


Figure 7: Dice score (top) and 95HD (bottom) comparison of automated thalamic segmentations to manual delineations of 10 HCP subjects. Scores are stated for our previous structural only method as well as the three likelihood implementations of our joint method.

Dice comparison of each likelihood implementation on these two sets of reconstructed tensors can be seen in Fig. 8. These results generally show that all three models are reasonably robust to such an acquisition change in HCP quality data, with a median Dice score of 0.85 or greater in each grouped label across all models. Similarly, each model showed Dice scores of greater than 0.95 for whole thalamus outlines. While median Dice scores for each grouped label are comparable between models, the DSW-beta does appear to be more robust for whole thalamus as well as three of the five grouped labels. Additionally, both Wishart and Log-Gaussian show a larger number of low Dice outliers with minimum values reaching as low as 30 points below the median in some cases.

3.4. Applications to conventional quality dMRI

While our method assumes that the resolution of the diffusion MRI approaches 1 mm isotropic (which is the case for many modern datasets, e.g., following the HCP protocol), it is of high interest to segment the thalamic nuclei in lower resolution scans for two reasons. First, because large

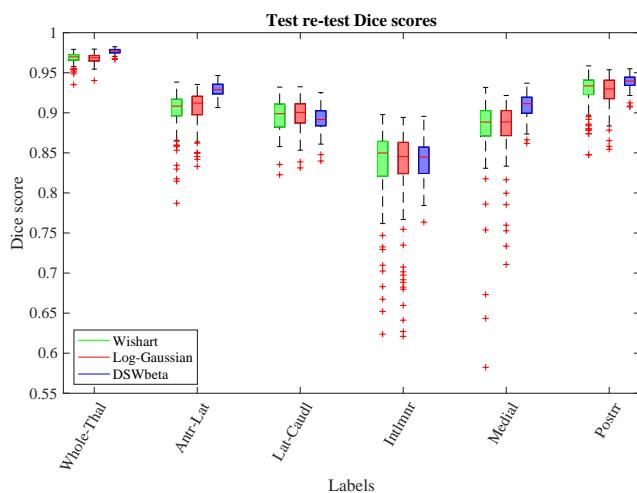


Figure 8: Dice score evaluation of test-retest reliability on 110 HCP subjects. For each subject, we performed two segmentations using DTI images obtained by fitting the tensor to the data from $b=1000$ s/mm^2 and $b=2000$ s/mm^2 shells separately and computed Dice scores for groups of labels in the two resulting segmentations.

amounts of legacy data were acquired at lower resolution. And second, because many current studies (e.g., ADNI, GENFI) still use those acquisitions, either in order not to deviate from the protocol used to acquire images earlier in the project or to accommodate acquisition constraints such as available scanner time. As explained in Section 2.3 above, compatibility with conventional quality data is actually the reason why we chose to model the diffusion tensor in our likelihood term, rather than using a more sophisticated, higher order model.

Reduced resolution and contrast on such scans, compared to HCP, make manual delineation using the criteria from Section 3.1 infeasible. For this reason we do not do a direct comparison of our methods to the 10 labels defined by manual segmentation. Instead, we first evaluate the reliability of the joint segmentation method through test-retest analysis, then assess the utility of our method, using the ability to discriminate individuals with AD vs controls as a proxy for accuracy.

3.4.1. Test-retest reliability analysis

In order to assess the test-retest reliability of the method on lower resolution dMRI, we used a separate dataset, comprising 21 healthy volunteers (9 male, 12 female, aged 53 – 80 years) acquired at the UCL DRC. Three MRI sequences were performed for each subject in a single session: one T1-weighted MPRAGE 1.1 mm isotropic resolution; and two diffusion weighted acquisitions each consisting of 64 gradient directions at a b -value of 1,000 s/mm^2 and a 2.5 mm isotropic resolution. Using the two dMRI acquisitions as separate tests, segmentations were performed at a 1 mm isotropic resolution in the native orientation of the individual dMRI volumes before being resampled to the native space of the structural volume for calculation of test-retest Dice

Multi-modal thalamic segmentation

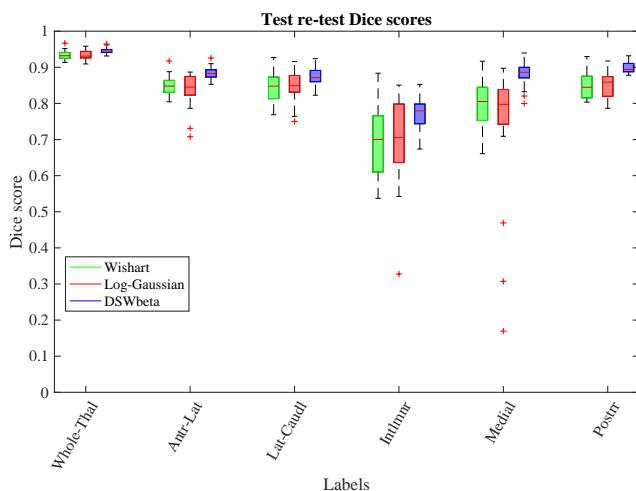


Figure 9: Dice score evaluation of test-retest reliability on conventional-quality data from 21 subjects acquired at the UCL Dementia Research Centre. For each subject, we performed two segmentations using dMRI data acquired in the same session using the same acquisition parameters and computed Dice scores for groups of labels.

1 scores. The Dice scores for this experiment are shown in
2 Fig. 9.

3 While this experiment shows lower Dice scores than in
4 HCP, possibly due to the increased voxel size and reduced
5 quality of the data, median scores are still above 0.9 for
6 whole thalamus and 0.8 for four of the five grouped labels.
7 However, it is more apparent from this plot that both the
8 Wishart and Log-Gaussian implementations are less robust
9 in lower quality data. This may be due to the increased
10 dimensionality of these two models, meaning imprecise fitting
11 of the tensor model caused by partial volume effects has a
12 greater impact than for the more robust FA and principle
13 direction model used by the DSW-beta likelihood.

3.4.2. Alzheimer's disease study

15 In order to evaluate the usefulness of our method in
16 a classical group study with imaging data of conventional
17 quality, we ran a further experiment using the ADNI dataset.
18 The ADNI was launched in 2003 by the National Institute
19 on Aging, the National Institute of Biomedical Imaging and
20 Bioengineering, the Food and Drug Administration, private
21 pharmaceutical companies and non-profit organizations, as
22 a \$60 million, 5-year public-private partnership. The main
23 goal of ADNI is to test whether MRI, positron emission
24 tomography (PET), other biological markers, and clinical
25 and neuropsychological assessment can be combined to
analyze the progression of MCI and early AD. Markers of
26 early AD progression can aid researchers and clinicians to
27 develop new treatments and monitor their effectiveness, as
28 well as decrease the time and cost of clinical trials. The
29 Principal Investigator of this initiative is Michael W. Weiner,
30 MD, VA Medical Center and University of California — San
31 Francisco. ADNI is a joint effort by co-investigators from
32 industry and academia. Subjects have been recruited from
33 over 50 sites across the U.S. and Canada. The initial goal
34

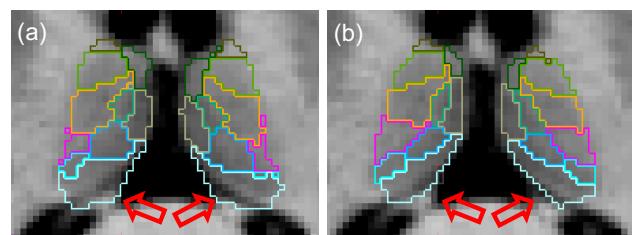


Figure 10: Comparison of thalamic segmentations of a subject from the ADNI dataset using equal (a) and reduced (b) dMRI likelihood weighting. Weighting the dMRI likelihood by the ratio of voxel volumes between sMRI and dMRI results in more accurate estimation of boundaries with heavy partial voluming in the diffusion channel, e.g., the CSF/posterior-thalamus boundary (red arrows).

of ADNI was to recruit 800 subjects but ADNI has been followed by ADNI-GO and ADNI-2. These three protocols have recruited over 1500 adults (ages 55–90) to participate in the study, consisting of cognitively normal older individuals, people with early or late MCI, and people with early AD. The follow up duration of each group is specified in the corresponding protocols for ADNI-1, ADNI-2 and ADNI-GO. Subjects originally recruited for ADNI-1 and ADNI-GO had the option to be followed in ADNI-2. For up-to-date information, see <http://www.adni-info.org>.

Specifically, we repeated the experiment from our previous article (Iglesias et al., 2019) using the T1-weighted and dMRI scans of 92 subjects from ADNI2. Here we considered 47 subjects with AD and 45 age-matched controls (73.8 ± 7.7 years; 44 females total). The data consisted of T1-weighted scans, with a resolution of $1.2 \times 1 \times 1$ mm (sagittal), and dMRI with a resolution of $1.35 \times 1.35 \times 2.7$ mm (axial). We fit the DTI model to the $b=1000$ s/mm² shell (41 directions), combined with 5 volumes at $b=0$. Given the test-retest results above and considering the similar levels of accuracy seen between dMRI likelihood models in the ground truth comparisons, we focus on our DSW-beta likelihood model in this experiment.

Initial tests on subjects from the ADNI dataset revealed some cases where the inclusion of the dMRI shifts boundaries in the segmentation due to the lower resolution of the dMRI data (and thus increased partial volume effects). An example is the over-segmentation of the thalamus into the CSF in Fig. 10a. We addressed this by allowing the contribution of the dMRI likelihood term to be reduced in proportion to the ratio between voxel volumes in the sMRI and dMRI volumes (Fig. 10b) as outlined in Section 2.5 and Section S.2 of the supplement.

As in Iglesias et al. 2019, we computed receiver operating characteristic (ROC) curves for discrimination of subjects into the two classes using five approaches: three based on thresholding the volume of the whole thalamus (as given by the FreeSurfer recon-all stream, the structural segmentation, and the joint segmentation); and two based on thresholding the likelihood ratio given by a linear discriminant analysis (LDA, Fisher 1936) on the volumes of the histological nuclei (as given by the structural and

Multi-modal thalamic segmentation

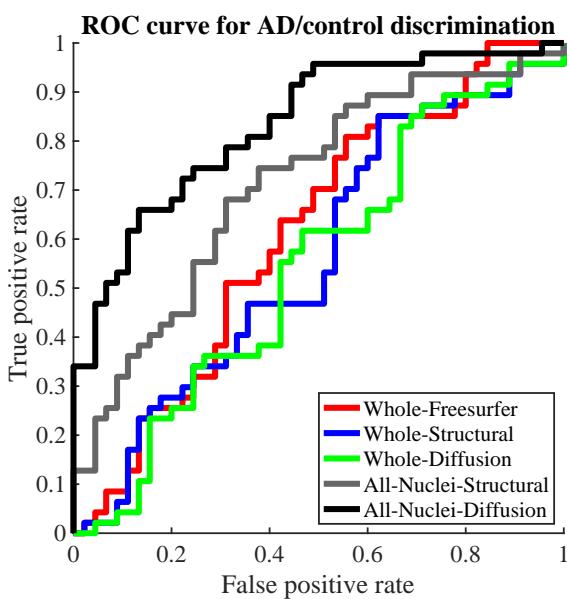


Figure 11: ROC curves for subjects with AD vs controls classification based on thalamic volumes.

Table 2

AUC, accuracy at elbow, and p-value for improved AUC values as given by a DeLong test.

	FreeSurfer (whole)	Structural (nuclei)	Diffusion (nuclei)
AUC	61.28%	71.30%	83.36%
Acc. at elbow	61.96%	68.48%	76.09%
p-value vs FreeSurfer		1.6e-01	1.1e-03
p-value vs Structural			2.1e-02

joint segmentation). The resulting ROC curve is shown in Fig. 11 with the area under the curve (AUC), accuracy at the elbow and p-values for comparison of AUC values shown in Table 2.

From these curves we can see that all three methods relying on the total volume of the thalamus have poor discriminative ability, with little difference between using FreeSurfer, structural or joint segmentations. This contrasts to the nuclei specific methods, which both show marked improvements. Structural showing an increase of 10% AUC over FreeSurfer's whole thalamus and joint segmentation an increase of 22%. However, only the improvements of the joint method show statistical significance with $p = 1.1e-03$ vs. FreeSurfer and $p = 2.1e-02$ vs structural nuclei segmentation.

4. Discussion and Conclusion

In this article, we have presented and tested a novel segmentation method for thalamic subregions from structural and diffusion MRI. Building on the Bayesian segmentation literature, we use novel likelihood models to exploit structural and diffusion MRI information *jointly* in order

to obtain an accurate parcellation of the thalamic nuclei. The information in structural MRI enables placement of boundaries in regions with strong contrast (e.g., medial boundary with the ventricles) with high precision, attributed to its higher resolution; the diffusion information enables the accurate segmentation of boundaries that are invisible in typical structural MRI sequences. Furthermore, we have presented an improved version of our previous histological atlas, which enables more accurate modelling of diffusion MRI in the cerebral white matter. The proposed method will be distributed with FreeSurfer and is widely applicable because the likelihood: (i) relies on a simple DTI model, which makes it compatible with virtually every diffusion dataset; (ii) adjusts to different resolutions by correcting for voxel sizes; and (iii) relies on an unsupervised model that is robust against changes in MR contrast.

We have conducted extensive experiments with manual segmentations, test-retest acquisition, and group studies – including datasets with different resolutions. The results have shown that the joint model exploiting the diffusion information improves accuracy over structural-only segmentation. Moreover, we have also found that the varying resolution gap between structural and diffusion MRI may be accommodated by weighting the diffusion likelihood term to account for voxel size differences, thus bypassing the need to explicitly model partial voluming – which quickly becomes intractable, particularly in multi-modal images defined on different voxel grids. While both our proposed likelihood model (DSW-beta) and the two competing alternatives showed similar levels of improved accuracy over structural-only segmentation, we found the DSW-beta distribution to have the highest test-retest reliability and to be the most robust to changes in the DTI resolution. Compared with other approaches, we produce Dice scores that are within an expected range. For example Su et al. (2019) reported mean scores of 0.64 and above, but direct comparison is complicated by differences in label definition, acquisition type and image resolution.

Our proposed method has a large number of design choices, particularly linked to the specification of shared parameters across classes in the structural and diffusion mixture models. We set these parameters with the combination of expert prior knowledge, a labelled template, and a well-known approach from the decision making literature (TOPSIS). While this approach is suboptimal (our prior knowledge is imperfect; a single template is biased towards a certain population, contrast, and resolution; and TOPSIS's criteria may not necessarily be ideal), it yielded groupings that worked well in practice for different datasets with different resolution.

This work has two main limitations. First, the lack of quantitative validation of our adapted manual segmentation against other segmentation criteria for the thalamus, e.g., with intra- and inter-rater variability. And second, the lack of explicit modelling for the partial volume effect; while accounting for the voxel size ratio mitigated this problem in our experiments, it is possible that it does not suffice

Multi-modal thalamic segmentation

1 for more extreme ratios. Addressing these two issues with
2 further experiments and solutions based on CNNs remain as
3 future work.

4 The presented method will be publicly available in
5 FreeSurfer as an extension of our current structural-only
6 code. As high-resolution diffusion data become increasingly
7 accessible, algorithms that can exploit them to produce
8 accurate segmentations – particularly for boundaries that are
9 invisible in structural MRI – have the potential to greatly
10 enhance neuroimaging studies.

11 Acknowledgments

12 This work was primarily funded by Alzheimer's Re-
13 search UK (ARUK-IRG2019A003). PG's work in this area
14 was supported by NIH NIBIB NAC P41EB015902 AY's
15 work in this area was supported by NIH grants R01 EB021265
16 and R56 MH121426. DCA's work in this area was supported
17 by EPSRC grant EP/R006032/1 and Wellcome Trust award
18 221915/Z/20/Z. The Dementia Research Centre is supported
19 by Alzheimer's Research UK, Alzheimer's Society, Brain
20 Research UK, and The Wolfson Foundation. This work was
21 supported by the National Institute for Health Research
22 (NIHR) Queen Square Dementia Biomedical Research Unit
23 and the University College London Hospitals Biomedical
24 Research Centre, the Leonard Wolfson Experimental Neu-
25 rology Centre (LWENC) Clinical Research Facility, and the
26 UK Dementia Research Institute, which receives its funding
27 from UK DRI Ltd, funded by the UK Medical Research
28 Council, Alzheimer's Society and Alzheimer's Research
29 UK. MB is supported by a Fellowship award from the
30 Alzheimer's Society, UK (AS-JF-19a-004-517). MB's work
31 was also supported by the UK Dementia Research Institute
32 which receives its funding from DRI Ltd, funded by the
33 UK Medical Research Council, Alzheimer's Society and
34 Alzheimer's Research UK. JDR is supported by the Miriam
35 Marks Brain Research UK Senior Fellowship and has re-
36 ceived funding from an MRC Clinician Scientist Fellowship
37 (MR/M008525/1) and the NIHR Rare Disease Translational
38 Research Collaboration (BRC149/NS/MH). JEI is supported
39 by the European Research Council (Starting Grant 677697,
40 project BUNGEETOOLS) and the NIH (1RF1MH123195-
41 01 and 1R01AG070988-01).

42 The collection and sharing of the ADNI data was funded
43 by the Alzheimer's Disease Neuroimaging Initiative (Na-
44 tional Institutes of Health Grant U01 AG024904) and De-
45 partment of Defence (W81XWH-12-2-0012). ADNI is funded
46 by the National Institute on Aging, the National Institute of
47 Biomedical Imaging and Bioengineering, and the follow-
48 ing: Alzheimer's Association; Alzheimer's Drug Discovery
49 Foundation; BioClinica, Inc.; Biogen Idec Inc.; Bristol-
50 Myers Squibb Company; Eisai Inc.; Elan Pharmaceuticals,
51 Inc.; Eli Lilly and Company; F. Hoffmann-La Roche Ltd and
52 affiliated company Genentech, Inc.; GE Healthcare; Inno-
53 genetics, N.V.; IXICO Ltd.; Janssen Alzheimer Immunother-
54 apy Research & Development, LLC.; Johnson & Johnson
55 Pharmaceutical Research & Development LLC.; Medpace,

1 Inc.; Merck & Co., Inc.; Meso Scale Diagnostics, LLC.;
2 NeuroRx Research; Novartis Pharmaceuticals Corporation;
3 Pfizer Inc.; Piramal Imaging; Servier; Synarc Inc.; and
4 Takeda Pharmaceutical Company. The Canadian Institutes
5 of Health Research is providing funds for ADNI clinical
6 sites in Canada. Private sector contributions are facilitated
7 by the Foundation for the National Institutes of Health. The
8 grantee is the Northern California Institute for Research and
9 Education, and the study is coordinated by the Alzheimer's
10 Disease Cooperative Study at the University of California,
11 San Diego. ADNI is disseminated by the Laboratory for
12 Neuro Imaging at the University of Southern California.

13 References

14 Al-Saady, M. L., Wolf, N. I., Pouwels, P. J. W., 2022. Segmentation of
15 intrinsically very low contrast magnetic resonance brain images using
16 tensor-based DTI registration. *Neuroimage: Reports* 2 (4), 100120.
17 Andersson, J. L. R., Sotiropoulos, S. N., Jan 2016. An integrated approach
18 to correction for off-resonance effects and subject movement in diffusion
19 MR imaging. *Neuroimage* 125, 1063–1078.
20 Aron, A. R., Schlaghecken, F., Fletcher, P. C., Bullmore, E. T., Eimer, M.,
21 et al., 2003. Inhibition of subliminally primed responses is mediated by
22 the caudate and thalamus: Evidence from functional MRI and Hunting-
23 ton's disease. *Brain* 126 (3), 713–723.
24 Arsigny, V., Fillard, P., Pennec, X., Ayache, N., 2006. Log-Euclidean
25 metrics for fast and simple calculus on diffusion tensors. *Magnetic
26 Resonance in Medicine* 56 (2), 411–421.
27 Ashburner, J., Andersson, J., Friston, K., 2000. Image registration using
28 a symmetric prior - in three dimensions. *Human Brain Mapping* 9 (4),
29 212–225.
30 Ashburner, J., Friston, K. J., 2005. Unified segmentation. *NeuroImage*
31 26 (3), 839 – 851.
32 Basile, G. A., Bertino, S., Bramanti, A., Ciurleo, R., Anastasi, G. P.,
33 et al., 6/14/2022 2021. In vivo super-resolution track-density imaging
34 for thalamic nuclei identification. *Cerebral Cortex* 31 (12), 5613–5636.
35 Battistella, G., Najdenovska, E., Maeder, P., Ghazaleh, N., Daducci, A. o.,
36 2017. Robust thalamic nuclei segmentation method based on local
37 diffusion magnetic resonance properties. *Brain Structure and Function*
38 222 (5), 2203–2216.
39 Behrens, T. E., Johansen-Berg, H., Woolrich, M., Smith, S., Wheeler-
40 Kingshott, C., et al., 2003. Non-invasive mapping of connections be-
41 between human thalamus and cortex using diffusion imaging. *Nature
42 neuroscience* 6 (7), 750–757.
43 Behzadian, M., Khammohammadi Otaghsara, S., Yazdani, M., Ignatius, J.,
44 2012. A state-of-the-art survey of TOPSIS applications. *Expert Systems
45 with Applications* 39 (17), 13051–13069.
46 Billot, B., Greve, D. N., Van Leemput, K., Fischl, B., Iglesias, J. E.,
47 et al., 06–08 Jul 2020. A learning strategy for contrast-agnostic MRI
48 segmentation. In: Arbel, T., Ben Ayed, I., de Bruijne, M., Descoteaux,
49 M., Lombaert, H., Pal, C. (Eds.), *Proceedings of the Third Conference
50 on Medical Imaging with Deep Learning*. Vol. 121 of *Proceedings of
51 Machine Learning Research*. PMLR, pp. 75–93.
52 Bocchetta, M., Gordon, E., Cardoso, M. J., Modat, M., Ourselin, S., et al.,
53 2018. Thalamic atrophy in frontotemporal dementia—not just a C9orf72
54 problem. *NeuroImage: Clinical* 18, 675–681.
55 Braak, H., Braak, E., 1991a. Alzheimer's disease affects limbic nuclei of
56 the thalamus. *Acta Neuropathol* 81 (3), 261–268.
57 Braak, H., Braak, E., 1991b. Neuropathological stageing of Alzheimer-
58 related changes. *Acta Neuropathol* 82 (4), 239–259.
59 Clayden, J. D., Thomas, D. L., Kraskov, A., 2019. Tractography-based
60 parcellation does not provide strong evidence of anatomical organisation
61 within the thalamus. *NeuroImage* 199, 418–426.
62 Dale, A. M., Fischl, B., Sereno, M. I., 1999. Cortical surface-based analysis:
63 I. segmentation and surface reconstruction. *NeuroImage* 9 (2), 179–194.

Multi-modal thalamic segmentation

1 de Jong, L. W., van der Hiele, K., Veer, I. M., Houwing, J. J., Westendorp, R.
2 G. J., et al., 6/14/2022 2008. Strongly reduced volumes of putamen and
3 thalamus in Alzheimer's disease: an MRI study. *Brain* 131 (12), 3277–
4 3285.

5 Dempster, A. P., Laird, N. M., Rubin, D. B., 1977. Maximum likelihood
6 from incomplete data via the EM algorithm. *Journal of the Royal
7 Statistical Society: Series B (Methodological)* 39 (1), 1–22.

8 Dryden, I. L., Koloydenko, A., Zhou, D., 2009. Non-Euclidean statistics for
9 covariance matrices, with applications to diffusion tensor imaging. *The
10 Annals of Applied Statistics* 3 (3), 1102 – 1123.

11 Eickhoff, S. B., Thirion, B., Varoquaux, G., Bzdok, D., Dec 2015.
12 Connectivity-based parcellation: Critique and implications. *Human
13 Brain Mapping* 36 (12), 4771–4792.

14 Fama, R., Sullivan, E. V., 2015. Thalamic structures and associated cognitive
15 functions: Relations with age and aging. *Neuroscience & Biobehavioral
16 Reviews* 54, 29–37.

17 Fan, Q., Witzel, T., Nummenmaa, A., Van Dijk, K. R., Van Horn, J. D.,
18 et al., 2016. MGH-USC human connectome project datasets with ultra-
19 high b-value diffusion MRI. *NeuroImage* 124, 1108–1114.

20 Fischl, B., 2012. FreeSurfer. *NeuroImage* 62 (2), 774–781.

21 Fischl, B., Salat, D. H., Busa, E., Albert, M., Dieterich, M., et al., 2002.
22 Whole brain segmentation: Automated labeling of neuroanatomical
23 structures in the human brain. *Neuron* 33 (3), 341–355.

24 Fischl, B., Sereno, M. I., Dale, A. M., 1999. Cortical surface-based analysis:
25 II: Inflation, flattening, and a surface-based coordinate system. *NeuroImage*
26 9 (2), 195–207.

27 Fischl, B., van der Kouwe, A., Destrieux, C., Halgren, E., Ségonne, F.,
28 et al., 01 2004. Automatically Parcellating the Human Cerebral Cortex.
29 *Cerebral Cortex* 14 (1), 11–22.

30 Fisher, R. A., 2022/06/23 1936. The use of multiple measurements in
31 taxonomic problems. *Annals of Eugenics* 7 (2), 179–188.

32 Görür, D., Rasmussen, C. E., Jul 2010. Dirichlet process gaussian mixture
33 models: Choice of the base distribution. *Journal of Computer Science
34 and Technology* 25 (4), 653–664.

35 Greve, D. N., Fischl, B., Oct 2009. Accurate and robust brain image
36 alignment using boundary-based registration. *Neuroimage* 48 (1), 63–
37 72.

38 Heckemann, R. A., Hajnal, J. V., Aljabar, P., Rueckert, D., Hammers, A. o.,
39 2006. Automatic anatomical brain MRI segmentation combining label
40 propagation and decision fusion. *NeuroImage* 33 (1), 115–126.

41 Henderson, J. M., Carpenter, K., Cartwright, H., Halliday, G. M., 6/14/2022
42 2000. Loss of thalamic intralaminar nuclei in progressive supranuclear
43 palsy and Parkinson's disease: clinical and therapeutic implications.
44 *Brain* 123 (7), 1410–1421.

45 Henschel, L., Conjeti, S., Estrada, S., Diers, K., Fischl, B., et al., 2020.
46 Fastsurfer - a fast and accurate deep learning based neuroimaging
47 pipeline. *NeuroImage* 219, 117012.

48 Hwang, K., Bertolero, M. A., Liu, W. B., D'Esposito, M., 06 2017. The
49 human thalamus is an integrative hub for functional brain networks. *The
50 Journal of Neuroscience* 37 (23), 5594.

51 Iglesias, J. E., Augustinack, J. C., Nguyen, K., Player, C. M., Player, A.,
52 et al., 2015. A computational atlas of the hippocampal formation using
53 ex vivo, ultra-high resolution MRI: Application to adaptive segmenta-
54 tion of in vivo MRI. *NeuroImage* 115, 117–137.

55 Iglesias, J. E., Insausti, R., Lerma-Usabiaga, G., Bocchetta, M., Leemput,
56 K. V., et al., 2018. A probabilistic atlas of the human thalamic nuclei
57 combining ex vivo MRI and histology. *NeuroImage* 183, 314 – 326.

58 Iglesias, J. E., Van Leemput, K., Golland, P., Yendiki, A., 2019. Joint in-
59 ference on structural and diffusion MRI for sequence-adaptive Bayesian
60 segmentation of thalamic nuclei with probabilistic atlases. In: Chung, A.
61 C. S., Gee, J. C., Yushkevich, P. A., Bao, S. (Eds.), *Information Process-
62 ing in Medical Imaging*. Springer International Publishing, Cham, pp.
63 767–779.

64 Jack Jr., C. R., Bernstein, M. A., Fox, N. C., Thompson, P., et al., 2022/06/28
65 2008. The Alzheimer's disease neuroimaging initiative (ADNI): MRI
66 methods. *Journal of Magnetic Resonance Imaging* 27 (4), 685–691.

67 Jakab, A., Blanc, R., Berényi, E. L., Székely, G., 12 2012. Generation of
68 individualized thalamus target maps by using statistical shape models
and thalamocortical tractography. *American Journal of Neuroradiology*
33 (11), 2110.

69 Jeurissen, B., Tournier, J.-D., Dhollander, T., Connelly, A., Sijbers, J., 2014.
70 Multi-tissue constrained spherical deconvolution for improved analysis
71 of multi-shell diffusion MRI data. *NeuroImage* 103, 411–426.

72 Jian, B., Vemuri, B. C., 2007. Multi-fiber reconstruction from diffusion MRI
73 using mixture of wisharts and sparse deconvolution. In: Karssemeijer,
74 N., Lelieveldt, B. (Eds.), *Information Processing in Medical Imaging*.
75 Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 384–395.

76 Johansen-Berg, H., Behrens, T., Sillery, E., Ciccarelli, O., Thompson, A.,
77 et al., 2005. Functional-anatomical validation and individual variation
78 of diffusion tractography-based segmentation of the human thalamus.
79 *Cerebral Cortex* 15 (1), 31–39.

80 Jones, E. G., 2012. The thalamus. Springer Science & Business Media.

81 Joshi, S., Davis, B., Jomier, M., Gerig, G., 2004. Unbiased diffeomorphic
82 atlas construction for computational anatomy. *NeuroImage* 23, S151–
83 S160.

84 Kassubek, J., Juengling, F. D., Ecker, D., Landwehrmeyer, G. B., 2005.
85 Thalamic atrophy in Huntington's disease co-varies with cognitive per-
86 formance: A morphometric MRI analysis. *Cerebral Cortex* 15 (6), 846–
87 853.

88 Krauth, A., Blanc, R., Poveda, A., Jeanmonod, D., et al., 2010. A mean
89 three-dimensional atlas of the human thalamus: Generation from multi-
90 ple histological data. *NeuroImage* 49 (3), 2053–2062.

91 Liu, Y., D'Haese, P.-F., Newton, A. T., Dawant, B. M., 2020. Generation of
92 human thalamus atlases from 7 T data and application to intrathalamic
93 nuclei segmentation in clinical 3 T T1-weighted images. *Magnetic
94 Resonance Imaging* 65, 114–128.

95 Maffei, C., Lee, C., Planich, M., Ramprasad, M., Ravi, N., et al., 2021.
96 Using diffusion MRI data acquired with ultra-high gradient strength to
97 improve tractography in routine-quality data. *NeuroImage* 245, 118706.

98 Mai, J. K., Forutan, F., 2012. Chapter 19 - Thalamus. Academic Press, San
99 Diego, pp. 618–677.

100 Mai, J. K., Majtanik, M., 2019. Toward a common terminology for the
101 thalamus. *Frontiers in Neuroanatomy* 12.

102 Mang, S. C., Busza, A., Reiterer, S., Grodd, W., Klose, U., 2022/06/14 2012.
103 Thalamus segmentation based on the local diffusion direction: A group
104 study. *Magnetic Resonance in Medicine* 67 (1), 118–126.

105 Mardia, K. V., Jupp, P. E., Mardia, K., 2000. *Directional statistics*. Vol. 2.
106 Wiley Online Library.

107 McKenna, M. C., Li Hi Shing, S., Murad, A., Lope, J., Hardiman, O., et al.,
108 2022. Focal thalamus pathology in frontotemporal dementia: Phenotype-
109 associated thalamic profiles. *Journal of the Neurological Sciences* 436,
110 120221.

111 Minagar, A., Barnett, M. H., Benedict, R. H. B., Pelletier, D., Pirko, I., et al.,
112 01 2013. The thalamus and multiple sclerosis. *Neurology* 80 (2), 210.

113 Morel, A., 2007. *Stereotactic atlas of the human thalamus and basal ganglia*.
114 CRC Press.

115 Patenaude, B., Smith, S. M., Kennedy, D. N., Jenkinson, M., 2011. A
116 Bayesian model of shape and appearance for subcortical brain segmen-
117 tation. *NeuroImage* 56 (3), 907–922.

118 Planche, V., Su, J. H., Mournet, S., Saranathan, M., et al., 2022/06/28
119 2019. White-matter-nulled MPRAGE at 7T reveals thalamic lesions
120 and atrophy of specific thalamic nuclei in multiple sclerosis. *Multiple
121 Sclerosis Journal* 26 (8), 987–992.

122 Pohl, K. M., Fisher, J., Grimson, W. E. L., Kikinis, R., Wells, W. M., 2006.
123 A Bayesian model for joint segmentation and registration. *NeuroImage*
124 31 (1), 228–239.

125 Puonti, O., Iglesias, J. E., Leemput, K. V., 2016. Fast and sequence-adaptive
126 whole-brain segmentation using parametric Bayesian modeling. *Neu-
127 roImage* 143, 235 – 249.

128 Rohrer, J. D., Nicholas, J. M., Cash, D. M., van Swieten, J., Doppler, E.,
129 et al., 2015. Presymptomatic cognitive and neuroanatomical changes in
130 genetic frontotemporal dementia in the genetic frontotemporal dementia
131 initiative (GENFI) study: a cross-sectional analysis. *The Lancet Neurol-
132 ogy* 14 (3), 253–262.

133 Sadikot, A. F., Mallar Chakravarty, M., Bertrand, G., Rymar, V. V.,
134 et al., 2011. Creation of computerized 3D MRI-integrated atlases of

Multi-modal thalamic segmentation

1 the human basal ganglia and thalamus. *Frontiers in Systems Neuro-*
2 *science (SEPTEMBER 2011).*

3 Schmahmann, J. D., 2022/06/14 2003. Vascular syndromes of the thalamus.
4 *Stroke* 34 (9), 2264–2278.

5 Semedo, C., Cardoso, M. J., Vos, S. B., Sudre, C. H., Bocchetta, M., et al.,
6 2018. Thalamic nuclei segmentation using tractography, population-
7 specific priors and local fibre orientation. In: Frangi, A. F., Schnabel,
8 J. A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (Eds.), *Medical*
9 *Image Computing and Computer Assisted Intervention –MICCAI 2018*.
10 Springer International Publishing, Cham, pp. 383–391.

11 Sherman, S. M., 2007. The thalamus is more than just a relay. *Current*
12 *Opinion in Neurobiology* 17 (4), 417–422.

13 Sherman, S. M., 2016. Thalamus plays a central role in ongoing cortical
14 functioning. *Nature Neuroscience* 19 (4), 533–541.

15 Sherman, S. M., Guillery, R. W., 2001. *Exploring the thalamus*. Elsevier.

16 Smith, S. M., Jenkinson, M., Woolrich, M. W., Beckmann, C. F., Behrens,
17 T. E. J., et al., 2004. Advances in functional and structural MR image
18 analysis and implementation as FSL. *NeuroImage* 23, S208–S219.

19 Su, J. H., Thomas, F. T., Kasoff, W. S., Tourdias, T., Choi, E. Y., et al.,
20 2019. Thalamus optimized multi atlas segmentation (THOMAS): fast,
21 fully automated segmentation of thalamic nuclei from structural MRI.
22 *Neuroimage* 194, 272–282.

23 Tax, C. M., Jeurissen, B., Vos, S. B., Viergever, M. A., et al., 2014. Recursive
24 calibration of the fiber response function for spherical deconvolution
25 of diffusion mri data. *NeuroImage* 86, 67–80.

26 Tourdias, T., Saranathan, M., Levesque, I. R., Su, J., Rutt, B. K., 2014. Vi-
27 sualization of intra-thalamic nuclei with optimized white-matter-nulled
28 MPRAGE at 7T. *NeuroImage* 84, 534–545.

29 Umapathy, L., Keerthivasan, M. B., Zahr, N. M., Bilgin, A., Saranathan, M.,
30 2021. Convolutional neural network based frameworks for fast automatic
31 segmentation of thalamic nuclei from native and synthesized contrast
32 structural MRI. *Neuroinformatics*.

33 Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E., Yacoub, E.,
34 et al., 2013. The WU-Minn human connectome project: An overview.
35 *NeuroImage* 80, 62–79.

36 Van Leemput, K., 2009. Encoding probabilistic brain atlases using Bayesian
37 inference. *IEEE Transactions on Medical Imaging* 28 (6), 822–837.

38 Van Leemput, K., Maes, F., Vandermeulen, D., Suetens, P., 1999. Auto-
39 mated model-based tissue classification of MR images of the brain. *IEEE*
40 *Transactions on Medical Imaging* 18 (10), 897–908.

41 Vatsavayai, S. C., Yoon, S. J., Gardner, R. C., Gendron, T. F., Vargas, J.
42 N. S., et al., Dec 2016. Timing and significance of pathological fea-
43 tures in C9orf72 expansion-associated frontotemporal dementia. *Brain*
44 139 (Pt 12), 3202–3216.

45 Wachinger, C., Reuter, M., Klein, T., 2018. DeepNAT: Deep convolutional
46 neural network for segmenting neuroanatomy. *NeuroImage* 170, 434–
47 445.

48 Wells, W. M., Grimson, W. E. L., Kikinis, R., Jolesz, F. A., 1996. Adaptive
49 segmentation of MRI data. *IEEE Transactions on Medical Imaging*
50 15 (4), 429–442.

51 Yendiki, A., Panneck, P., Srinivasan, P., Stevens, A., Zöllei, L., et al.,
52 2011. Automated probabilistic reconstruction of white-matter pathways
53 in health and disease using an atlas of the underlying anatomy. *Frontiers*
54 *in Neuroinformatics* 5.

55 Zarei, M., Patenaude, B., Damoiseaux, J., Morgese, C., Smith, S., et al.,
56 2010. Combining shape and connectivity analysis: An MRI study of
57 thalamic degeneration in Alzheimer's disease. *NeuroImage* 49 (1), 1–8.

58 Zhang, D., Snyder, A. Z., Fox, M. D., Sansbury, M. W., Shimony, J. S., et al.,
59 2008. Intrinsic functional relations between human cerebral cortex and
60 thalamus. *Journal of neurophysiology* 100 (4), 1740–1748.

61 Zhang, H., Schneider, T., Wheeler-Kingshott, C. A., Alexander, D. C., 2012.
62 NODDI: Practical in vivo neurite orientation dispersion and density
63 imaging of the human brain. *NeuroImage* 61 (4), 1000–1016.

64 Zhang, Y., Brady, M., Smith, S., 2001. Segmentation of brain MR images
65 through a hidden Markov random field model and the expectation-
66 maximization algorithm. *IEEE Transactions on Medical Imaging* 20 (1),
67 45–57.