Cytoarchitecture Measurements in Brain Gray Matter using Likelihood-Free Inference

Maëliss Jallais¹^[0000–0001–5939–388X],
Pedro L. C. Rodrigues¹^[0000–0003–3534–317X],
Alexandre Gramfort¹^[0000–0001–9791–4404], and
Demian Wassermann¹^[0000–0001–5194–6056]

Université Paris-Saclay, Inria, CEA, Palaiseau, 91120, France
`firstname.lastname@inria.fr`

Abstract. Effective characterisation of the brain grey matter cytoarchitecture with quantitative sensitivity to soma density and volume remains an unsolved challenge in diffusion MRI (dMRI). Solving the problem of relating the dMRI signal with cytoarchitectural characteristics calls for the definition of a mathematical model that describes brain tissue via a handful of physiologically-relevant parameters and an algorithm for inverting the model. To address this issue, we propose a new forward model, specifically a new system of equations, requiring six relatively sparse b-shells. These requirements are a drastic reduction of those used in current proposals to estimate grey matter cytoarchitecture. We then apply current tools from Bayesian analysis known as likelihood-free inference (LFI) to invert our proposed model. As opposed to other approaches from the literature, our LFI-based algorithm yields not only an estimation of the parameter vector θ that best describes a given observed data point $\mathbf{x}_o$, but also a full posterior distribution $p(\theta|\mathbf{x}_o)$ over the parameter space. This enables a richer description of the model inversion results providing indicators such as confidence intervals for the estimations, and better understanding of the parameter regions where the model may present indeterminacies. We approximate the posterior distribution using deep neural density estimators, known as normalizing flows, and fit them using a set of repeated simulations from the forward model. We validate our approach on simulations using `dmipy` and then apply the whole pipeline to the HCP MGH dataset.

Keywords: Diffusion MRI · Brain Microstructure · Likelihood-Free Inference.

1 Introduction

Obtaining quantitative measurements of brain grey matter microstructure with a dedicated soma representation is a growing field of interest in the diffusion MRI (dMRI) community [18, 13]. However, current methods require demanding acquisitions and stabilise parameter fitting by enforcing constraints which are not valid across all brain tissue. Yet, these still encounter large indetermination

areas in the solution space making the results unstable [12]. To overcome these limitations, we propose a new model based on diffusion signal summary measurements, computed from boundary approximations. Our model has the benefit of reducing acquisition requirements while retaining microstructure sensitivity. Combining it with a likelihood-free inference (LFI) method from Bayesian analysis, we readily compute the posterior distribution of the model fitting at each voxel. Real asset to this method, it describes the quality of the fit. Notably, it provides confidence intervals and better understanding of the parameter regions where the model may present indeterminacies.

Accessing cortical cytoarchitecture *in vivo* is a sought step to understand diseases such as Alzheimer's. Unlike histology, diffusion MRI enables to quantify brain tissue characteristics non-invasively. Our method could, for example, help understanding dementia and cognitive deficits, which appear to be mostly related to neuronal loss and synaptic pathology. Relationships between cellular microstructure and cognition could also be established [9]. A major asset brought by the proposed method is its reliance on only a few relatively sparse q -shells, equivalently b -shells, along with confidence intervals that help to guarantee the good fitting of the parameters.

Current microstructure models are predominantly based on the two compartment Standard Model (SM) [18, 11]. Recent evidence shows that the SM, mainly used in white matter, does not hold for grey matter microstructure analysis [17]. Several assumptions aim at explaining this issue such as increased permeability in neurite membranes [17], or curvy projections along with longer pulse duration [11]. We follow the hypothesis that the SM doesn't hold due to an abundance of cell bodies in gray matter [13]. Our proposed biophysical model is then based on three compartments [13]: neurites; somas; and extra-cellular space (ECS). Despite its increased complexity, the main advantage of such model is the possibility to jointly estimate the characteristic features of each compartment.

By proposing a forward model for grey matter using three-compartments, we are able to simulate the dMRI signal for known tissue parameters. However, using it to quantify microstructure from the signal remains tricky. Several approaches have tried to tackle this non-linear inverse problem. Still, most popular solutions to date, such as NODDI [e.g. 18], stabilise this inverse problem by imposing constraints on model parameters which are not biologically plausible [12]. This biases parameter estimation and the inverse problem remains largely degenerate. Our proposal tackles these limitations based on three contributions. First, we introduce a new parameter that jointly encodes soma radius and inner diffusivity without imposing constraints on these values. This new parameter reduces indeterminacies in the model and has relevant physiological interpretations. Second, we present a new method to fit the model through summary features of the dMRI signal based on a large and small q -value analysis. These rotationally-invariant features relate directly to the tissue parameters, and enable us to invert the model without manipulating the raw dMRI signals. These summary statistics allow to have a stable solution of the parameter estimations, as opposed to the indeterminate models used in [18, 13]. Third, we employ modern tools

from Bayesian analysis known as likelihood-free inference (LFI) to determine the posterior distribution of the fitted parameters. This posterior offers a full description of the solution landscape and can point out degeneracies.

The likelihood function of complex simulators is often very hard to obtain and, therefore, determining the posterior distribution via Bayes’s formula is hopeless. LFI bypasses this bottleneck by recurring to several simulations of the model using different parameters and learning the posterior distribution from these examples. The first works on LFI are also known as approximate Bayesian computation (ABC) and have been applied to invert models from ecology, population genetics, and epidemiology [16]. Recently, there has been a growing interest in the machine learning community in improving the limitations of ABC methods. These limitations include the large number of simulations required for the posterior estimations or the need of defining a distance function to compare the results of two simulations. A special class of neural networks called normalizing flows has been used to approximate the posterior distribution of simulator models and has demonstrated promising results in different research fields [8].

In what follows, we present our three contributions and validate our workflow for inverting dMRI signals on a set of physiologically relevant situations using the `dmipy` simulator. Then, we apply our method on the HCP MGH dataset, which contains in-vivo human brain acquisitions, and present the results of parameter estimations in grey matter along with an assessment of their variances.

2 Methods

2.1 Modeling the Brain Gray Matter with a 3-compartment Model

To characterize cortical cytoarchitecture, we propose a method to relate the diffusion MRI signal to specific tissue parameters. To that aim, we first need to define a model of the grey matter tissue. Histology teaches us that grey matter is composed of neurons embedded in a fluid environment. Each neuron is composed of a soma, corresponding to the cellular body, surrounded by neurites connecting neurons together. Following this tissue biophysical composition, we model the grey matter tissue as three-compartmental [13], moving away from the SM designed for white matter. We are assuming that our acquisition protocol is not sensitive to exchanges between those compartments, i.e. molecules moving from one compartment to another have a negligible influence on the signal [13]. The acquired signal is then considered as resulting from a convex mixture of signals arising from somas, neurites, and extra-cellular space (ECS). Unlike white matter-centric methods, we are not interested in the fiber orientation and only estimate orientation-independent parameters. This enables us to work on the direction-averaged dMRI signal, denoted $\bar{S}(q)$, known as the powder averaged signal. This consideration mainly matters for neurites, as their signal is not isotropic, as opposed to the proposed model for somas and ECS. Our direction-averaged grey matter signal model is then:

$$\frac{\bar{S}(q)}{\bar{S}(0)} = f_n \bar{S}_{\text{neurites}}(q, D_a) + f_s \bar{S}_{\text{somas}}(q, D_s, r_s) + f_{\text{ECS}} \bar{S}_{\text{ECS}}(q, D_e). \quad (1)$$

In this equation, f_n ; f_s ; and f_{ECS} represent signal fractions for neurites; somas; and ECS respectively ($f_n + f_s + f_{ECS} = 1$). Note that the relative signal fractions do not correspond to the relative volume fractions of the tissue compartments as they are also modulated by a difference in T2 values [11]. D_a corresponds to axial diffusivity inside neurites, while D_s and D_e correspond to somas and extra-cellular diffusivities. r_s is the average soma radius within a voxel. We use q -values for more readability and harmonization throughout the paper, but a direct conversion to b -values is also possible, using $b = (2\pi q)^2 \tau$ with $\tau = \Delta - \delta/3$.

We now review the model for each compartment, to make explicit the impact of each parameter on the diffusion MRI signal.

Neurite compartment. Neurites, as in the SM, are modeled as 0-radius impermeable cylinders (“sticks”), with effective diffusion along the parallel axis, and a negligible radial intra-neurite diffusivity. In our acquisition setting, this model has been shown to hold [17]. Its direction averaged signal is [17]:

$$\bar{S}_{\text{neurites}}(q) \simeq \frac{1}{4\sqrt{\pi\tau D_a}} \cdot q^{-1} \quad (2)$$

Soma compartment. Somas are modeled as spheres, whose signal can be computed using the GPD approximation [3]:

$$-\log \bar{S}_{\text{somas}}(q) = C(r_s, D_s) \cdot q^2 \quad (3)$$

We exploit this relation here to extract a parameter $C_s = C(r_s, D_s)[m^2]$ which, at fixed diffusivity D_s , is modulated by the radius of the soma r_s :

$$C(r_s, D_s) = \frac{2}{D_s \delta^2} \sum_{m=1}^{\infty} \frac{\alpha_m^{-4}}{\alpha_m^2 r_s^2 - 2} \cdot \left(2\delta - \frac{2 + e^{-\alpha_m^2 D_s (\Delta - \delta)} - e^{-\alpha_m^2 D_s \delta} - e^{-\alpha_m^2 D_s \Delta} + e^{-\alpha_m^2 D_s (\Delta + \delta)}}{\alpha_m^2 D_s} \right)$$

where α_m is the m th root of $(\alpha r_s)^{-1} J_{\frac{3}{2}}(\alpha r_s) = J_{\frac{5}{2}}(\alpha r_s)$, with $J_n(x)$ the Bessel functions of the first kind.

Extra-cellular space compartment. The extra-cellular space is approximated as isotropic Gaussian diffusion, i.e. a mono-exponential diffusion signal with a scalar diffusion constant D_e , which reflects the molecular viscosity of the fluid. This approximation assumes that the ECS is fully connected. The approximation is therefore:

$$-\log(\bar{S}_{\text{ECS}}(q)) = (2\pi q)^2 \tau D_e \quad (4)$$

We are estimating D_e from the ventricles, given the same metabolic composition of the extracellular space and ventricles: cerebrospinal fluid.

2.2 An Invertible 3-compartment Model: dMRI Summary Statistics

The tissue model presented in Section 2.1 enables to relate the dMRI signal with parameters representing grey matter tissue microstructure. However, solving the

inverse problem directly from Eq. (1) is a difficult task, leading to indeterminacies and bad parameter estimations. Current methods addressing this issue have not studied its stability [13] but simpler models with only two compartments have been shown to be indeterminate [11].

To produce a method which addresses this indeterminacy, we introduce rotationally invariant summary statistics to describe the dMRI signal. Furthermore, we solve the inverse problem efficiently reducing indeterminacies and we develop a method to detect them in Section 2.3. These dMRI-based summary statistics are extracted from our proposed model presented in Section 2.1 via the following analysis of the dMRI signal on the boundaries of large and small q -value cases.

Large q -value approximation: RTOP. We compute a q -bounded RTOP, a direct measure of the restrictions of the diffusing fluid molecule motion that gives us information about the structure of the media [10]:

$$\text{RTOP}(q) = 4\pi \int_0^q \frac{\bar{S}(\eta)}{S(0)} \eta^2 d\eta \quad (5)$$

For q large enough, RTOP on our 3-compartment model in Eq. (1), yields a soma and extra-cellular signal which converges towards a constant value in q , while the neurites' contribution becomes quadratic in q . In this case, RTOP becomes:

$$\text{RTOP}(q) = \underbrace{f_s \left(\frac{\pi}{C_s} \right)^{3/2}}_{a_{\text{fit}}} + \frac{f_{\text{ECS}}}{8(\pi\tau D_e)^{3/2}} + \underbrace{\frac{f_n}{2} \cdot \sqrt{\frac{\pi}{\tau D_a}}}_{b_{\text{fit}}} \cdot q^2 \quad (6)$$

By accurately estimating the second derivative of $\text{RTOP}(q)$ at q large enough, we can solve the coefficients of the polynomial in Eq. (6): a_{fit} and b_{fit} . We do this efficiently by casting it as an overdetermined ordinary least squares problem which has a unique solution.

Small q -value approximation: Spiked LEMONADE. We propose a second approximation, based on a moment decomposition for small q -values [12]:

$$\frac{S_{\hat{\mathbf{g}}}(q)}{S(0)} = 1 - b(q)M_{i_1 i_2}^{(2)} g_1 g_2 + \frac{b(q)^2}{2!} M_{i_1 \dots i_4}^{(4)} g_1 \dots g_4 - \dots, \quad b(q) = (2\pi q)^2 \tau \quad (7)$$

where i_k are the directional basis of the tensors M , $g_k = i_k \cdot \hat{\mathbf{g}} \in \mathbb{R}^3$, and $\hat{\mathbf{g}}$ the unit direction of the dMRI acquisition. From the moment tensors of this decomposition, LEMONADE [12] extracts rotational invariant scalar indices $M^{(i),j}$, $i, j \in \{0, 2, 4, \dots\}$. These quantify white matter microstructure by plugging the 2-compartment SM into Eq. (7) [see 12, app. C].

In this work, we extended LEMONADE to our 3-compartment model presented in Section 2.1 by plugging Eq. (1) into Eq. (7) and performing tedious arithmetic. This results in the following equation system, which now includes the

soma parameter C_s , relating the dMRI signal to gray matter microstructure:

$$\begin{cases} M^{(2),0} = f_n D_a + 3f_s \frac{C_s}{(2\pi)^2 \tau} + 3f_{\text{ECS}} D_e \\ M^{(2),2} = f_n D_a p_2 \\ M^{(4),0} = f_n D_a^2 + 5f_s \left(\frac{C_s}{(2\pi)^2 \tau} \right)^2 + 5f_{\text{ECS}} D_e^2 \\ M^{(4),2} = f_n D_a^2 p_2 \end{cases} \quad (8)$$

where p_2 is a scalar measure of neurite orientation dispersion [12].

Note that only the shells with $b(q) \leq 3 \text{ ms } \mu\text{m}^{-2}$ are used, to get an unbiased estimation of the rotational invariant moments $M^{(2),0}$, $M^{(2),2}$, $M^{(4),0}$ and $M^{(4),2}$.

Complete system. Combining equations (6) and (8) and adding the constraint that the fractions for the three compartments sum to one, we obtain a system of 7 equations and 7 unknowns. Following Menon et al. [9], we assume D_e nearly-constant per subject acquisition and that it can be estimated as the mean diffusivity in the subject's ventricles. This assumption allows us to drop an unknown from the system, use D_e as a reference diffusivity and render our variables unitless as $D_a^u = \frac{D_a}{D_e}$ and $C_s^u = \frac{C_s}{(2\pi)^2 \tau D_e}$, obtaining:

Spiked LEMONADE	RTOP
Small q -values	Large q -values
$\begin{cases} \frac{M^{(2),0}}{D_e^2} = f_n D_a^u + 3f_s C_s^u + 3f_{\text{ECS}} \\ \frac{M^{(2),2}}{D_e^2} = f_n D_a^u \cdot p_2 \\ \frac{M^{(4),0}}{D_e^2} = f_n D_a^{u2} + 5f_s C_s^{u2} + 5f_{\text{ECS}} \\ \frac{M^{(4),2}}{D_e^2} = f_n D_a^{u2} \cdot p_2 \end{cases}$	$\begin{cases} a_{\text{fit}} (\tau D_e)^{3/2} = \frac{f_s}{8(\pi C_s^u)^{3/2}} + \frac{f_{\text{ECS}}}{8\pi^{3/2}} \\ b_{\text{fit}} (\tau D_e)^{1/2} = \frac{f_n}{2} \sqrt{\frac{\pi}{D_a^u}} \end{cases}$
$f_n + f_s + f_{\text{ECS}} = 1$	

2.3 Solving the inverse problem via likelihood free inference

The Bayesian formalism. The system of equations presented in Section 2.2 describes how the tissue parameters $\boldsymbol{\theta} = (D_a^u, C_s^u, p_2, f_s, f_n, f_{\text{ECS}}) \in \mathbb{R}^6$ are related to the vector of summary statistics obtained from the raw dMRI signal

$$\mathbf{x} = \left(\frac{M^{(2),0}}{D_e^2}, \frac{M^{(2),2}}{D_e^2}, \frac{M^{(4),0}}{D_e^2}, \frac{M^{(4),2}}{D_e^2}, a_{\text{fit}} (\tau D_e)^{3/2}, b_{\text{fit}} \sqrt{\tau D_e} \right) \in \mathbb{R}^7.$$

We denote this relation by $\mathcal{M} : \mathbb{R}^6 \rightarrow \mathbb{R}^7$ such that $\mathbf{x} = \mathcal{M}(\boldsymbol{\theta})$. We interpret the inverse problem of inferring the parameters $\boldsymbol{\theta}_o$ that best describe a given observed summary feature vector $\mathbf{x}_o$ as that of determining the posterior distribution of $\boldsymbol{\theta}$ given an observation $\mathbf{x}_o$. By first choosing a prior distribution $p(\boldsymbol{\theta})$ describing our initial knowledge of the parameter values, we may use Bayes' theorem to write

$$p(\boldsymbol{\theta}|\mathbf{x}_o) = \frac{p(\mathbf{x}_o|\boldsymbol{\theta})p(\boldsymbol{\theta})}{p(\mathbf{x}_o)}, \quad (9)$$

where $p(\mathbf{x}_o|\boldsymbol{\theta})$ is the likelihood of the observed data point and $p(\mathbf{x}_o)$ is a normalizing constant. Note that such probabilistic approach returns not only which

θ best fits the observed data (i.e. the parameter that maximizes the posterior distribution) but also describes the variance of the parameters being estimated. Furthermore, it provides a full description of which regions of the parameter space may generate the same vector of observed summary features.

Bypassing the likelihood function. Despite its apparent simplicity, it is usually difficult to use Eq. (9) to determine the posterior distribution, since the likelihood function for data points generated by complex non-linear models is often hard to write. To avoid such difficulty, we directly approximate the posterior distribution using a class of neural networks called normalizing flows [14]. These flows are invertible functions capable of transforming vectors generated by a simple base distribution (e.g. the standard multivariate Gaussian distribution) into an approximation of the true posterior distribution. We denote a normalizing flow by $q_\phi^N(\theta|\mathbf{x})$, where ϕ parametrizes the neural network, N indicates how many data points were available when training the neural density estimator, and θ and $\mathbf{x}$ are input arguments. In this work, we use normalizing flows with an autoregressive architecture implemented via the masked autoencoder for distribution estimation (MADE) [6]. We follow the same setup from [8] for LFI problems, stacking five MADEs, each with two hidden layers of 50 units, and a standard normal base distribution. This choice provides a sufficiently flexible function capable of approximating complex posterior distributions. We refer the reader to [14] for more information on the different types of normalizing flows.

Our setup for approximating $p(\theta|\mathbf{x}_o)$ with a normalizing flow is the following: first, generate a set of N paired samples $(\theta_i, \mathbf{x}_i)$, where $\theta_i \sim p(\theta)$ and $\mathbf{x}_i = \mathcal{M}(\theta_i)$. The parameters ϕ are then obtained via stochastic gradient descent to minimize the loss function

$$\mathcal{L}(\phi) = - \sum_{i=1}^N \log \left(q_\phi^N(\theta_i|\mathbf{x}_i) \right). \quad (10)$$

Note that (10) is a Monte-Carlo approximation to the Kullback-Leibler (KL) divergence between $q_\phi^N(\theta|\mathbf{x})$ and $p(\theta|\mathbf{x})$, so the parameter ϕ that minimizes it yields a normalizing flow which is the closest to the true posterior distribution in the KL-sense. Moreover, it is possible to show that the minimizer of Eq. (10) converges to $p(\theta|\mathbf{y})$ when $N \rightarrow \infty$ [8]. We obtain the posterior distribution for observation $\mathbf{x}_o$ by simply writing $p(\theta|\mathbf{x}_o) \approx q_\phi^N(\theta|\mathbf{x}_o)$.

3 Results and discussion

3.1 Simulations

Validating LFI on a base case. In this experiment, our goal was to validate the use of the LFI procedure described in Section 2.3 to approximate the posterior distribution of the tissue parameters in a simple setting, where their ground truth values are known. We report the results with a single choice of ground truth parameters, but our actual complete validation was performed on a set of physiologically relevant choices of parameters.

Fig. 1: Histograms of 10^4 samples of the approximate posterior distribution with observed dMRI signals generated under two acquisition setups, $\mathcal{A}$ and $\mathcal{B}$ (see text for details). Vertical black dashed lines represent ground truth values of θ_o which generated the observed signals. Different colors show how the posterior distribution gets sharper as the number N of simulations in the training dataset increases. Solid curves indicate results for setup $\mathcal{A}$, which are very close to the true values, and dashed curves for setup $\mathcal{B}$, which present a bias.

Experimental setup. We simulated the three-compartment model described in Section 2.1 using the `dmipy` simulator [5]. Somas have been modeled as spheres with a $35 \mu\text{m}$ diameter and $2.3 \mu\text{m}^2 \text{ms}^{-1}$ diffusivity; neurites as tubes with $1 \mu\text{m}$ diameter and $1.7 \mu\text{m}^2 \text{ms}^{-1}$ axial diffusivity; and extra-cellular space as isotropic Gaussian diffusion with a diffusivity of $2.6 \mu\text{m}^2 \text{ms}^{-1}$. These values have been selected to mimic existing neurons available from the NeuroMorpho database (neuromorpho.org). In what follows, we denote this choice of tissue parameters θ_o . We have considered two acquisition setups for generating diffusion signals, with $\delta = 12.9 \text{ms}$ and $\Delta = 21.8 \text{ms}$. These values correspond to the acquisition parameters used in the HCP MGH dataset. Both setups have b -shells with 128 uniformly distributed directions, but they differ in their b -values. Setup $\mathcal{A}$ corresponds to an “ideal” case with 10 b -values between 0 and $10 \text{ms} \mu\text{m}^{-2}$. Setup $\mathcal{B}$ reproduces the more challenging setup from the HCP dataset, with only 5 b -values: 0, 1, 3, 5 and $10 \text{ms} \mu\text{m}^{-2}$. We simulated a diffusion signal for each setup and computed their summary statistics, obtaining $\mathbf{x}_o^{\mathcal{A}}$ and $\mathbf{x}_o^{\mathcal{B}}$. Note that in setup $\mathcal{B}$ we have added an extra step when extracting the summary features: we used MAPL [4] to interpolate an additional b -value to the simulations at $0.1 \text{ms} \mu\text{m}^{-2}$, aiming to improve the close-to-zero approximation of the moments (as presented in Section 2.2), which requires 3 b -values inferior to $2.5 \text{ms} \mu\text{m}^{-2}$.

We trained a neural density estimator $q_\phi^N(\theta|\mathbf{x})$ using the procedure described in Section 2.3 with N simulations from the three-compartment model and an uniform prior distribution defined on physiologically relevant intervals:

$$D_a \in [10^{-5}, 3.5], C_s \in [50, 2500], \text{ and } p_2, f_s, f_n, f_{\text{ECS}} \in [0, 1]. \quad (11)$$

By the end of the optimization procedure that minimizes Eq. (10), we obtain approximations $q_\phi^N(\theta|\mathbf{x}_o^{\mathcal{A}}) \approx p(\theta|\mathbf{x}_o^{\mathcal{A}})$ and $q_\phi^N(\theta|\mathbf{x}_o^{\mathcal{B}}) \approx p(\theta|\mathbf{x}_o^{\mathcal{B}})$.

The estimated posterior distribution. Figure 1 portrays the marginal posterior distributions for each tissue parameter in setup $\mathcal{A}$ with $N \in \{10^3, 10^4, 10^5\}$ and setup $\mathcal{B}$ with $N = 10^5$. We see that the marginals of $q_\phi^N(\theta|\mathbf{x}_o^{\mathcal{A}})$ become

Fig. 2: (A) C_s dependence on soma radius r_s and diffusivity D_s . We see that there are several values of (d_s, D_s) that yield the same C_s . (B) Histograms of 10^4 samples from the marginal and joint posterior distributions of $d_s = 2r_s$ and D_s . The ridge in the joint distribution indicates that there are several possible values for the pair (d_s, D_s) with high probability, which are those yielding the same C_s . Estimating C_s directly bypasses this indeterminacy.

sharper around the ground truth values θ_o when N increases. This indicates that the LFI method is able to correctly inverse the proposed model in Eq. (1) based on the summary statistics presented in Section 2.1 when enough b -values are available. The posterior marginals of $q_\phi^N(\theta|\mathbf{x}_o^B)$ present a bias with respect to the ground truth parameters, most likely due to the reduced number of b -values available in this setup. Note, however, that the bias is relatively low for the C_s estimation, indicating that our procedure can still be useful for describing the somas sizes in the HCP MGH dataset, as we show in Section 3.2.

C_s avoids model indeterminacy. In Section 2.1, we introduce the parameter C_s , which serves as a proxy of the soma radius and provides key information on the soma compartment. Estimating this parameter instead of the soma radius (r_s) and diffusivity (D_s) separately avoids falling into model indeterminacy problems, such as those observed in the literature [13]. Indeed, Figure 2 presents the marginal posterior distributions of r_s and D_s as well as their joint distribution in the $\mathcal{A}$ setup; we extend the prior distribution intervals from Eq. 11 to include $r_s \in [10^{-5}, 30]$ and $D_s \in [1, 3]$ with the ground truth parameters $r_s = 17.5$ and $D_s = 2.3$. In addition to larger marginal posterior distributions for each parameter, the joint posterior shows a valley of values for the (r_s, D_s) pair, including the ground truth parameters, with high probability. In this situation the LFI procedure is unable to determine among all the possible solutions which one is the ground truth. Using C_s instead of r_s and D_s separately, we avoid such indeterminacy, as seen in Figure 1. Note that such behavior derives naturally from the way we define our model and inference procedure. As such, similar results can be expected when applying our method to the HCP dataset.

Assessing the quality of the posterior estimations. Deriving the posterior distributions of the parameter vectors allows us to report the values of the most likely tissue parameters for a given observation, along with our certitude regarding our inference. Figure 3 presents the logarithm of the standard deviation of the marginal posterior samples for different ground truth parameter choices (varying f_s and f_n) under setup $\mathcal{A}$. These values indicate how sharp a posterior

Fig. 3: Logarithm of the standard deviations for the marginal posterior distribution of D_a , C_s , and p_2 with different choices of ground truth parameters (varying f_s and f_n). We see that when the signal fraction of somas decreases ($f_s \rightarrow 0$) the standard deviation of the C_s -estimation increases; and when less neurites are present ($f_n \rightarrow 0$) the standard deviation of p_2 and D_a increase.

distribution is and, therefore, quantify the quality of the fit. We observe larger standard deviations in the absence (or weak presence) of soma compartments in the mixture signal, e.g., the standard deviation of C_s is large when few or no somas are present ($f_s \approx 0$). This is to be expected, the lack of contribution from the somas in diffusion signal making it difficult to estimate parameters related to them. Conversely, a low standard deviation is observed for signals generated in grey matter tissue conditions, where a soma predominance is expected.

3.2 HCP MGH results

We used the open data set HCP MGH Adult Diffusion [15] to study the variation of all the parameters on the human brain grey matter. It is composed of 35 subjects with $\delta/\Delta = 12.9/21.8$ ms and $b = 1, 3, 5, 10$ ms μm^{-2} . We used the 3, 5 and 10 ms μm^{-2} b -values for the RTOP approximation. For the Spiked LEMONADE approximation, we fitted MAPL on the 0, 1 and 3 ms μm^{-2} b -values and interpolated a point at 0.1 ms μm^{-2} to reduce noise and improve the estimation. D_e was estimated as the mean diffusivity in the ventricles. Our results are shown in Fig. 4. Overlay colormaps are masked showing only areas where parameters are stable. We deemed parameters stable when their value was larger than 2 times the LFI-obtained standard deviations of the fitted posterior. This results in a lack of stability on small sections including the auditory cortex and the precentral gyrus fundus. Our figure assesses qualitatively the results on soma size by comparing with nissl-stained histological studies [1, 2, 7]. This comparison shows good agreement between different cortical areas and our parameter C_s which, under nearly-constant intra-soma diffusion D_s , is modulated by soma size.

4 Conclusion

Quantifying grey matter tissue composition is challenging. In this work, we have presented a methodology to estimate the parameters of a model that best fit an

Fig. 4: Microstructural measurements averaged over 31 HCP MGH subjects. We deemed stable measurements with a z-score larger than 2, where the standard deviation on the posterior estimates was estimated through our LFI fitting approach. In comparing with Nissl-stained cytoarchitectural studies we can qualitatively evaluate our parameter C_s : Broadmann area 44 (A) has smaller soma size in average than area 45 (B) [2]; large von Economo neurons predominate the superior anterior insula (C) [1]; precentral gyrus (E) shows very small somas while post-central (D) larger ones [7].

observed data point, and also their full posterior distribution. This rich description provides many useful tools, such as assessing the quality of the parameter estimation or characterizing regions in the parameter space where it is harder to invert the model. Moreover, our proposal alleviates limitations from current methods in the literature by not requiring physiologically unrealistic constraints on the parameters and avoiding indeterminacies when estimating them.

Acknowledgement

This work was supported by the ERC-StG NeuroLang ID:757672 and the ANR BrAIN (ANR-20-CHIA-0016) grants.

References

1. Allman, J.M., Tetreault, N.A., Hakeem, A.Y., Manaye, K.F., Semendeferi, K., Erwin, J.M., Park, S., Goubert, V., Hof, P.R.: The von Economo neurons in fronto-insular and anterior cingulate cortex in great apes and humans. *Brain Structure and Function* **214**, 495–517 (Jun 2010)
2. Amunts, K., Schleicher, A., Bürgel, U., Mohlberg, H., Uylings, H.B., Zilles, K.: Broca’s region revisited: Cytoarchitecture and intersubject variability. *Journal of Comparative Neurology* **412**(2), 319–341 (1999)

3. Balinov, B., Jönsson, Linse, P., Söderman, O.: The NMR Self-Diffusion Method Applied to Restricted Diffusion. Simulation of Echo Attenuation from Molecules in Spheres and between Planes. *Journal of Magnetic Resonance, Series A* **104**(1), 17–25 (1993)
4. Fick, R.H., Wassermann, D., Caruyer, E., Deriche, R.: MAPL: Tissue microstructure estimation using Laplacian-regularized MAP-MRI and its application to HCP data. *NeuroImage* **134**, 365–385 (Jul 2016)
5. Fick, R.H.J., Wassermann, D., Deriche, R.: The Dmipy Toolbox: Diffusion MRI Multi-Compartment Modeling and Microstructure Recovery Made Easy. *Frontiers in Neuroinformatics* **13**, 64 (Oct 2019)
6. Germain, M., Gregor, K., Murray, I., Larochelle, H.: Made: Masked autoencoder for distribution estimation. In: *Proceedings of the 32nd International Conference on Machine Learning*, vol. 37, pp. 881–889, PMLR (2015)
7. Geyer, S., Schleicher, A., Zilles, K.: Areas 3a, 3b, and 1 of Human Primary Somatosensory Cortex. *NeuroImage* **10**(1), 63–83 (Jul 1999), ISSN 10538119
8. Greenberg, D., Nonnenmacher, M., Macke, J.: Automatic posterior transformation for likelihood-free inference. In: *Proceedings of the 36th International Conference on Machine Learning*, vol. 97, pp. 2404–2414, PMLR (2019)
9. Menon, V., Gallardo, G., Pinski, M.A., Nguyen, V.D., Li, J.R., Cai, W., Wassermann, D.: Microstructural organization of human insula is linked to its macrofunctional circuitry and predicts cognitive control. *eLife* **9**, e53470 (2020)
10. Mitra, P.P., Latour, L.L., Kleinberg, R.L., Sotak, C.H.: Pulsed-field-gradient NMR measurements of restricted diffusion and the return-to-origin probability. *Journal of Magnetic Resonance* **114**, 47–58 (1995)
11. Novikov, D.S., Fieremans, E., Jespersen, S.N., Kiselev, V.G.: Quantifying brain microstructure with diffusion MRI: Theory and parameter estimation: Brain microstructure with dMRI: Theory and parameter estimation. *NMR in Biomedicine* p. e3998 (Oct 2018), ISSN 09523480
12. Novikov, D.S., Veraart, J., Jelescu, I.O., Fieremans, E.: Rotationally-invariant mapping of scalar and orientational metrics of neuronal microstructure with diffusion MRI. *NeuroImage* **174**, 518–538 (Jul 2018)
13. Palombo, M., Ianus, A., Guerreri, M., Nunes, D., Alexander, D.C., Shemesh, N., Zhang, H.: SANDI: A compartment-based model for non-invasive apparent soma and neurite imaging by diffusion MRI. *NeuroImage* **215**, 116835 (2020)
14. Papamakarios, G., Nalisnick, E.T., Rezende, D.J., Mohamed, S., Lakshminarayanan, B.: Normalizing flows for probabilistic modeling and inference. *ArXiv abs/1912.02762* (2019)
15. Setsompop, K., Kimmlingen, R., Eberlein, E., Witzel, T., Cohen-Adad, J., McNab, J., Keil, B., Tisdall, M., Hoecht, P., Dietz, P., Cauley, S., Tountcheva, V., Matschl, V., Lenz, V., Heberlein, K., Potthast, A., Thein, H., Horn, J.V., Toga, A., Schmitt, F., Lehne, D., Rosen, B., Wedeen, V., Wald, L.: Pushing the limits of in vivo diffusion MRI for the Human Connectome Project. *NeuroImage* **80**, 220 – 233 (2013)
16. Sisson, S.A.: *Handbook of Approximate Bayesian Computation*. Chapman and Hall/CRC (Sep 2018)
17. Veraart, J., Nunes, D., Rudrapatna, U., Fieremans, E., Jones, D.K., Novikov, D.S., Shemesh, N.: Noninvasive quantification of axon radii using diffusion MRI. *eLife* **9** (Feb 2020)
18. Zhang, H., Schneider, T., Wheeler-Kingshott, C.A., Alexander, D.C.: NODDI: Practical in vivo neurite orientation dispersion and density imaging of the human brain. *NeuroImage* **61**(4), 1000 – 1016 (2012)