Published in final edited form as:

Neuroimage. 2018 July 15; 175: 272–285. doi:10.1016/j.neuroimage.2018.03.059.

Bayesian uncertainty quantification in linear models for diffusion MRI

Jens Sjölund^{a,b,c,*}, Anders Eklund^{b,c,d}, Evren Özarslan^{b,c}, Magnus Herberthson^e, Maria Bånkestad^f, and Hans Knutsson^{b,c}

^aElekta Instrument AB, Kungstensgatan 18, Box 7593, SE-103 93 Stockholm, Sweden

^bDepartment of Biomedical Engineering, Linköping University, Linköping, Sweden

^cCenter for Medical Image Science and Visualization (CMIV), Linköping University, Sweden

^dDepartment of Computer and Information Science, Linköping University, Linköping, Sweden

^eDepartment of Mathematics, Linköping University, Linköping, Sweden

^fRISE SICS, Isafjordsgatan 22, Box 1263, SE-164 29 Kista, Sweden

Abstract

Diffusion MRI (dMRI) is a valuable tool in the assessment of tissue microstructure. By fitting a model to the dMRI signal it is possible to derive various quantitative features. Several of the most popular dMRI signal models are expansions in an appropriately chosen basis, where the coefficients are determined using some variation of least-squares. However, such approaches lack any notion of uncertainty, which could be valuable in e.g. group analyses. In this work, we use a probabilistic interpretation of linear least-squares methods to recast popular dMRI models as Bayesian ones. This makes it possible to quantify the uncertainty of any derived quantity. In particular, for quantities that are affine functions of the coefficients, the posterior distribution can be expressed in closed-form. We simulated measurements from single- and double-tensor models where the correct values of several quantities are known, to validate that the theoretically derived quantiles agree with those observed empirically. We included results from residual bootstrap for comparison and found good agreement. The validation employed several different models: Diffusion Tensor Imaging (DTI), Mean Apparent Propagator MRI (MAP-MRI) and Constrained Spherical Deconvolution (CSD). We also used in vivo data to visualize maps of quantitative features and corresponding uncertainties, and to show how our approach can be used in a group analysis to downweight subjects with high uncertainty. In summary, we convert successful linear models for dMRI signal estimation to probabilistic models, capable of accurate uncertainty quantification.

Keywords

Diffusion MRI; uncertainty quantification; signal estimation

*Corresponding author: jens.sjoland@liu.se (Jens Sjölund).

Conflict of interest statement

Declarations of interest: none.

1. Introduction

Diffusion magnetic resonance imaging (dMRI) permits the noninvasive assessment of tissue microstructure. By fitting a model of the dMRI signal in each voxel it is possible to derive various quantitative features. Examples include voxel-based scalar indices such as fractional anisotropy (FA) or return to origin probability (RTOP), as well as the pairwise probability of a virtual fiber being traced between two points. Using such measures it is possible to perform statistical group analyses. For the reliability of such tests it is essential to quantify the uncertainty of the relevant measures. This fact is well-established in the field of functional magnetic resonance imaging (fMRI) — where it is common to, for example, downweight subjects with a high variance (Chen et al., 2012; Woolrich et al., 2004) — but not so for dMRI; the most popular approach for doing FA group comparisons (Smith et al., 2006) ignores uncertainty in FA.

A large body of research has been devoted to formulating new diffusion models that address the inability of diffusion tensor imaging (DTI) (Basser et al., 1994) to resolve crossing and kissing fibers. A common trait among several of the most widely used methods is that they expand the signal in an appropriately chosen functional basis. Often, the coefficients of the expansion are then determined by some variation of linear least-squares (see Table 1 for a non-exhaustive list). This is precisely the type of models we are concerned with in this paper — linear (in the coefficients) models fitted with least-squares. Incidentally, some non-parametric models (Andersson and Sotiropoulos, 2015; Sjölund et al., 2017) also belong to this class, although we will not elaborate on the connection here. Our key observation is that linear models fitted with least-squares are amenable to a probabilistic reinterpretation. We will show that it follows almost immediately that, under the same assumptions as in the fitting, it is possible to determine the full posterior distribution of the coefficients — not just a point estimate.

This is, of course, not the first time someone has taken a probabilistic view on signal estimation in dMRI. It is well-known that the noisy signal in MRI follows a Rician, or more generally a non-central Chi, distribution (Gudbjartsson and Patz, 1995), which on the other hand is approximately Gaussian when the signal-to-noise ratio is, at least, moderately high ($\gtrsim 3$). In DTI (Basser et al., 1994) and diffusion kurtosis imaging (DKI) (Jensen et al., 2005), it is common to fit a linear model to the logarithm of the signal. The resulting log-Rician distribution is again approximately normal for moderate signal-to-noise ratios, but with signal-dependent (heteroscedastic) noise (Salvador et al., 2005). In such cases, weighted least-squares has been shown to work well (Rawlings et al., 2001; Veraart et al., 2013). Error propagation in DTI fitted with nonlinear least-squares has also been investigated (Koay et al., 2007).

Bayesian methods (Behrens et al., 2003; Gelman et al., 2013; Gu et al., 2017; Wegmann et al., 2017) are distinctly different from methods using least-squares to find point estimates. By assuming parametric probability distributions for the likelihood and for every parameter in the model, Bayesian methods make it possible to derive the probability distribution of any

quantity of interest, at least in principle. Actually evaluating such distributions, however, typically relies extensively on sampling methods such as Markov Chain Monte Carlo.

Bootstrapping is a frequentistic alternative to Bayesian models. The general idea is to approximate the underlying probability distribution with an empirical one. Samples are drawn with replacement and for each draw the parameter of interest is calculated. By repeating this procedure, it is possible to approximate the sampling distribution of the relevant parameter. Naturally, the quality of the approximation degrades as the number of empirical samples decreases. Residual bootstrap (Chung et al., 2006) and wild bootstrap (Whitcher et al., 2008) are two forms of model-based bootstrapping that have been applied to dMRI. In residual bootstrap, the normalized residuals (after fitting a model) are resampled, and the fitting procedure is repeated for this new draw, after which the parameter value of interest is recorded. All residuals are assumed to have identical distributions and resampling is done freely among them. In wild bootstrap, on the other hand, modified residuals are randomly added or subtracted to the fitted point where they originated from, without being distributed to other design points. Model-based bootstrapping is only reliable insofar as the model can adequately describe the measured diffusion signals (Yuan et al., 2008). Basically all of the methods described above have been applied to tractography (Behrens et al., 2007; Berman et al., 2008; Haroon et al., 2009; Jeurissen et al., 2011; Jones, 2003, 2008).

In this work, instead of starting out by assuming more or less contrived prior distributions for the coefficients and the likelihood, we look at methods that are tried and tested and see what the corresponding priors are. Surprisingly, it results in a simple closed form expression for the posterior distribution of coefficients, and by extension also of all parameters linear in the coefficients. In other cases there is no need to repeat the whole fitting procedure, as in bootstrapping methods. Since the posterior is available in closed form it is very efficient to sample from it directly.

2. Theory

To recapitulate, we focus on linear models fitted with linear least-squares. This might sound restrictive at first but — as can be seen in Table 1 — it encompasses several of the most widely used models in dMRI. A broad distinction can be made between those derived from a cumulant expansion of the signal (Basser et al., 1994; Salvador et al., 2005; Veraart et al., 2013; Westin et al., 2016) and those where the signal is expanded in a, typically orthogonal, basis (Assemlal et al., 2009; Caruyer and Deriche, 2012; Descoteaux et al., 2006, 2007; Fick et al., 2016; Hess et al., 2006; Özarslan et al., 2013; Tournier et al., 2007; Tuch, 2004). Differences aside, all these models assume that, in the absence of noise, it would hold that

$$y(\mathbf{x}) = \sum_{i=1}^d c_i \phi_i(\mathbf{x}), \quad (1)$$

where y is the response variable that we want to model, $\mathbf{x}$ defines a measurement and $\phi_i(\mathbf{x})$ are (possibly nonlinear) functions with corresponding coefficients c_i . In a typical scenario,

we are given n observations $(\mathbf{x}_j, y_j)$ and the objective is to determine an estimate $\hat{\mathbf{c}}$ such that $y \approx \Phi \hat{\mathbf{c}}$, where we have introduced the matrix $\Phi_{ji} = \phi_i(\mathbf{x}_j)$. Most of the works cited in Table 1 simply state their estimate in one of these forms:

$$\hat{\mathbf{c}} = \begin{cases} (\Phi^T \Phi)^{-1} \Phi^T \mathbf{y} & \text{(ordinary least-squares)} \\ (\Phi^T W \Phi)^{-1} \Phi^T W \mathbf{y} & \text{(weighted least-squares)} \\ (\Phi^T \Phi + \Lambda)^{-1} \Phi^T \mathbf{y} & \text{(Tikhonov regularization).} \end{cases} \quad (2)$$

These are all special cases of the estimate

$$\hat{\mathbf{c}} = (\Phi^T W \Phi + \Lambda)^{-1} \Phi^T W \mathbf{y}, \quad (3)$$

but what is the probabilistic interpretation? In the next section, we show that equation (3) arises naturally as the maximum-a-posteriori solution of a Bayesian linear regression problem.

2.1. Least-squares and Bayesian linear regression

Suppose that the observations are given by adding Gaussian noise ϵ to the model in equation (1),

$$\mathbf{y}(\mathbf{x}) = \Phi \mathbf{c} + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \sigma^2 W^{-1}). \quad (4)$$

We assume that the matrix W , determining the structure of the noise, is known but not the scale σ^2 . It follows that the likelihood — the probability distribution of the observations given the parameters — is

$$p(\mathbf{y}|\mathbf{c}, \mathbf{x}) = \mathcal{N}(\Phi \mathbf{c}, \sigma^2 W^{-1}). \quad (5)$$

As is typical in the Bayesian formalism, we may specify a prior probability density over the coefficients $\mathbf{c}$. Suppose that we use

$$p(\mathbf{c}) = \mathcal{N}(0, \zeta^2 \Lambda^{-1}), \quad (6)$$

where Λ is known but not the scale ζ . We may assume that the expected value $\mu_{\mathbf{c}}$ is zero without loss of generality, since it is always possible to subtract $\Phi \mu_{\mathbf{c}}$ from the response variable. Bayes' rule then gives the posterior — the probability distribution of the coefficients given the data — as follows

$$p(\mathbf{c}|\mathbf{y}, \mathbf{x}) = \frac{p(\mathbf{y}|\mathbf{c}, \mathbf{x})p(\mathbf{c})}{p(\mathbf{y}|\mathbf{x})}. \quad (7)$$

Inserting the expressions for the likelihood and prior from above, it is a straight- forward exercise in probability theory (Bishop, 2006) to show that the posterior distribution is again a normal distribution, given by

$$\begin{aligned} p(\mathbf{c}|\mathbf{y}, \mathbf{x}) &= \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}), \\ \boldsymbol{\mu} &= \mathbf{Q}^{-1} \boldsymbol{\Phi}^T \mathbf{W} \mathbf{y}, \quad (8) \\ \boldsymbol{\Sigma} &= \sigma^2 \mathbf{Q}^{-1}, \end{aligned}$$

where $\mathbf{Q} = \boldsymbol{\Phi}^T \mathbf{W} \boldsymbol{\Phi} + (\sigma/\zeta)^2 \boldsymbol{\Lambda}$. We find that by setting $\zeta = \sigma$, the mean of the posterior distribution (8) becomes equal to the least-squares estimate (3). Thus, by changing perspective and considering the estimate (3) as the mean of the posterior distribution (8), we get—almost for free—a quantification of the uncertainty in the prediction. The only unknown variable that we need to estimate is the residual variance, σ^2 .

2.2. Residual variance

Before deriving an estimate for the residual variance, we first introduce the smoother matrix H , which maps responses to predictions, i.e. $\hat{\mathbf{y}} = \boldsymbol{\Phi} \hat{\mathbf{c}} = H\mathbf{y}$ (Hastie et al., 2001). Then, the covariance of the residuals, $\mathbf{r} = \mathbf{y} - \hat{\mathbf{y}}$, can be written

$$\text{Cov}[\mathbf{r}] = \text{Cov}[(\mathbf{I} - H)\mathbf{y}] \quad (9)$$

$$= (\mathbf{I} - H)\text{Cov}[\mathbf{y}](\mathbf{I} - H)^T \quad (10)$$

$$= (\mathbf{I} - H)\sigma^2 \mathbf{W}^{-1}(\mathbf{I} - H)^T \quad (11)$$

$$= \sigma^2 \mathbf{Z} \mathbf{Z}^T \quad (12)$$

where $\mathbf{Z} = (\mathbf{I} - H)\mathbf{W}^{-\frac{1}{2}}$ and $\mathbf{W}^{-\frac{1}{2}}$ denotes a matrix square root such as the one given by the Cholesky decomposition.

For an ordinary linear regression model with n observations and d basis functions, the residual variance is estimated as

$$\hat{\sigma}^2 = \frac{\|\mathbf{r}\|_2^2}{\nu}, \quad (13)$$

where $\nu = n - d$ is referred to as the degrees of freedom. We generalize the estimate of the residual variance to encompass both weighted- and regularized least-squares, by taking the trace of both sides in equation (12). This gives

$$\hat{\sigma}^2 = \frac{\|\mathbf{y} - \hat{\mathbf{y}}\|_2^2}{\|Z\|_F^2}. \quad (14)$$

In other words, the generalized expression for the the degrees of freedom is

$$\nu = \|Z\|_F^2 = \text{Tr}[(I - H)W^{-1}(I - H)^T]. \quad (15)$$

2.2.1. Residual bootstrap—In the results, section 3.1, we compare the uncertainty quantification that follows from the Bayesian interpretation in section 2.1 with that given by residual bootstrap (Chung et al., 2006).

Residual bootstrap is a form of model-based resampling (Davison and Hinkley, 1997), that has been successfully used for dMRI (Berman et al., 2008; Chung et al., 2006; Haroon et al., 2009; Jeurissen et al., 2011). In residual bootstrap, the normalized residuals (after fitting a model) are resampled, and the fitting procedure is repeated for this new draw, after which the parameter value of interest is recorded. The underlying assumption is that the normalized residuals have identical distributions, so that resampling can be done freely among them. For this to hold true, we see from equation (12) that the appropriate normalization is given by

$$\tilde{\mathbf{r}}_i = \frac{\mathbf{r}_i}{\sqrt{(ZZ^T)_{ii}}} \quad (16)$$

If $\Lambda = 0$, one can show by direct computation that $ZZ^T = (I - H)W^{-1}$. If, further (as assumed e.g. in DTI), the noise is independent but heteroscedastic, then W is diagonal but not constant. The expression for the modified residual (16) then simplifies to (Chung et al., 2006; Davison and Hinkley, 1997)

$$\tilde{\mathbf{r}}_i = \frac{\mathbf{r}_i}{\sqrt{(1 - H_{ii})W_{ii}^{-1}}}. \quad (17)$$

In other words, equation (16) actually extends the applicability of residual bootstrap, from the ordinary- or weighted least-squares methods previously described in the literature to cases with correlated noise and/or regularization.

2.3. Uncertainty of the residual variance estimate

It would be possible to insert the point estimate (14) of the residual variance in equation (8) to get a closed-form expression for the posterior distribution. However, this neglects the uncertainty in the residual variance estimation. We can handle this uncertainty in a Bayesian way, by specifying a distribution over the variance σ^2 . We then determine the posterior distribution for the coefficients by integrating over $p(\sigma^2)$; this procedure is referred to as marginalization.

In Appendix A we show that there is a particularly convenient prior (inverse-Gamma distribution) for which the posterior is a multivariate t -distribution (Kotz and Nadarajah, 2004; Roth, 2012),

$$p(\mathbf{c}|\mathbf{y}, \mathbf{x}) = t_{\nu}(\boldsymbol{\mu}, R), \quad (18)$$

where the mean $\boldsymbol{\mu}$, the degrees of freedom ν and the correlation matrix R are given by

$$\begin{aligned} \boldsymbol{\mu} &= Q^{-1} \Phi^T W \mathbf{y}, \\ \nu &= \|Z\|_F^2, \\ R &= \frac{\nu - 2}{\nu} \hat{\sigma}^2 Q^{-1}. \end{aligned} \quad (19)$$

We emphasize that this replaces equation (8) as the main expression for the uncertainty in the coefficients. Note that the mean, covariance and degrees of freedom are the same as before, but the t -distribution has fatter tails than the Gaussian distribution, reflecting the uncertainty in the variance. As the degrees of freedom increases it becomes more similar to a Gaussian distribution; a rule of thumb is that for $\nu \gtrsim 30$ the t -distribution can be replaced by a Gaussian.

2.4. Corollaries

We can imagine a multitude of uses for the posterior distribution (18).

- The multivariate t -distribution is closed under affine transformations (Roth, 2012), which means any property $\boldsymbol{\theta} = A\mathbf{c} + \mathbf{b}$ also follows a multivariate t -distribution, namely

$$\boldsymbol{\theta} \sim t_{\nu}(A\boldsymbol{\mu} + \mathbf{b}, ARA^T). \quad (20)$$

This makes it possible to, for instance, compute credible intervals (confidence intervals).

- The distribution of a property $\theta = f(\mathbf{c})$, nonlinearly dependent on the coefficients, can be determined analytically in some cases, notably if f is invertible. Regardless, it is always possible to estimate the distribution by sampling from the posterior.
- If the coefficients are used to condition a probability distribution $p(\theta|\mathbf{c})$, for instance when sampling fiber orientation in probabilistic tractography, one can, at least formally, write down the expression for the posterior predictive distribution:

$$p(\theta|\mathbf{y}, \mathbf{x}) = \int p(\theta|\mathbf{c})p(\mathbf{c}|\mathbf{y}, \mathbf{x})d\mathbf{c}. \quad (21)$$

2.5. Evaluation

We base most of our quantitative evaluation on the seemingly self-evident statement that “in p % of the cases the correct value should be smaller than the p th quantile”. To formalize this statement, recall the definition of the cumulative distribution function of a random variable X : $F_X(x) = \Pr(X \leq x)$. The quantile function Q is, under mild assumptions, the inverse cumulative distribution function, $Q = F^{-1}$. Stated differently, the quantile function specifies, for a given probability p , the value x at which the probability of the random variable is less than or equal to the given probability, $\Pr(X \leq x) = p$.

In our case, given a measurement $\mathbf{y}_j$ we get a posterior distribution of the quantity of interest θ_j following the line of reasoning presented in the previous sections. For a fixed probability p the corresponding quantile, $q_j = Q(p)$, can be estimated. If we know the correct value θ^* , we can observe empirically how often $\theta^* \leq q_j$ and compare this with the theoretically expected probability p . This procedure can be repeated for a range of probabilities p , after which the comparison can be presented graphically in a so called P-P (probability-probability) plot (Wilk and Gnanadesikan, 1968).

We also use the interquartile range (*IQR*) to summarize the dispersion in a single number. The interquartile range is defined as the difference between the third and the first quartile, i.e. $IQR = Q(0.75) - Q(0.25)$. For a symmetric distribution it would thus correspond to the 50% confidence interval. One of its main advantages is that it is more robust to outliers than, for instance, the standard deviation.

3. Results

We implemented the core algorithmic ideas by extending the free and open source software project Dipy (Diffusion Imaging in Python) (Garyfallidis et al., 2014). The code for our modified version of Dipy is available online¹. We made a quantitative analysis of the proposed method on simulated data, and a qualitative analysis on in vivo data from the Human Connectome Project (Van Essen et al., 2013).

¹<https://github.com/jsjol/dipy>

3.1. Simulated data

We simulated data from single- and double-tensor models. For these models, we can calculate the correct values of several quantities analytically.

In the single-tensor case the mean diffusivity (MD) and fractional anisotropy (FA) are given, respectively, by

$$MD = \frac{\text{Tr}(D)}{3}, \quad (22)$$

$$FA = \sqrt{\frac{1}{2} \left(3 - \frac{\text{Tr}(D)^2}{\text{Tr}(D^2)} \right)}, \quad (23)$$

where D is the diffusion tensor in question. The above, non-standard, expression for FA elucidates its nonlinear dependence on D (Özarslan et al., 2005).

In the double-tensor case, the crossing angle is known by design and can be compared with those returned by peak detection algorithms. In addition, there are some scalar indices that can be calculated analytically (Özarslan et al., 2013), e.g. the return to origin probability (RTOP)

$$RTOP = \frac{1}{S_0} \int_{\mathbb{R}^3} S(\mathbf{q}) d\mathbf{q}. \quad (24)$$

3.1.1. Single tensor model—Similarly to Chung et al. (2006), we simulated signals from three different Gaussian distributions with axially symmetric diffusion tensors, each having identical mean diffusivity, $MD = 0.7 \cdot 10^{-3} \text{ mm}^2/\text{s}$, but different fractional anisotropy, $FA \in \{0.2, 0.5, 0.8\}$. Assuming a Stejskal-Tanner type measurement (Stejskal and Tanner, 1965), the latent signal was

$$S(\mathbf{q}) = S_0 e^{-4\pi^2 t_d \mathbf{q}^T D \mathbf{q}}, \quad (25)$$

where $t_d = -\delta/3$; here t_d is the mixing time and δ the pulse duration. We used the same experimental parameters as in the Human Connectome data described in section 3.2.1, but with maximum b -value ($b = 4\pi^2 t_d q^2$) equal to 1000 s/mm^2 .

Based on the latent signal, we simulated 1000 independent measurements following a Rician distribution with scale $\sigma = 0.05 S_0$. For each measurement we used weighted least-squares to fit a diffusion tensor model. The coefficients in the fitting thus include D_{xx} , D_{yy} and D_{zz} . Recalling the definition of mean diffusivity (23), we see that it is a linear function of the

coefficients. Using equation (20) we find that the posterior distribution follows a univariate t -distribution. The P-P plot in Figure 1 shows the frequency with which the quantiles based on noisy measurements contain the correct value of the mean diffusivity and compares this with 1000 draws using residual bootstrap.

Fractional anisotropy, on the other hand, is a nonlinear function of the coefficients, which means that we have to resort to sampling to estimate its posterior distribution. For each measurement the quantiles were estimated from 1000 samples from the posterior of the coefficients. The resulting P-P plots are shown in Figure 2 and compared with 1000 draws using residual bootstrap.

3.1.2. Double tensor model—We defined the double tensor signal as the sum of two, equal magnitude, single tensor signals, but where the second was rotated by either 45° or 60° about the y -axis. The single tensor signal had mean diffusivity $0.7 \cdot 10^{-3} \text{ mm}^2/\text{s}$ and fractional anisotropy 0.8, which is roughly equivalent to a white matter tract (Pierpaoli and Basser, 1996). Again, we used the same experimental parameters as in the Human Connectome data described in section 3.2.1, but this time with maximum b -value equal to 3000 s/mm^2 . Based on the double tensor signal, we simulated 1000 independent measurements following a Rician distribution with scale $\sigma = 0.05 S_0$.

The return to origin probability (RTOP) for a multi tensor model with equal diffusion tensors is the same as for a single tensor model, namely

$$RTOP = \det(4\pi t_d D)^{-1/2}, \quad (26)$$

which in this case evaluates to $0.90 \cdot 10^6 \text{ mm}^{-3}$.

For each measurement we fitted a MAP-MRI model (Özarslan et al., 2013) with a Laplacian regularization determined using generalized cross-validation (Fick et al., 2016). For this model the estimate of the RTOP is a linear function of the coefficients, thus equation (20) applies. We found that the RTOP was on average overestimated. As shown in Figure 3, the size of the bias (difference between average and correct value) depends on both the Rician character of the signal and the maximum b -value used. With the settings used in our simulation (in particular, Rician measurements), the bias is smallest when $b = 3000 \text{ s/mm}^2$ and the maximum order of the MAP-MRI coefficients is 4. Nevertheless — even when using these settings — the bias has an immediate effect on a P-P plot as evident from Figure 4. We did, however, find that it was possible to remove this confounding effect from the P-P plot by subtracting the bias from the theoretical quantiles before comparing with the correct value, see Figures 4b and 4d. This corresponds to shifting the mean of the posterior distribution to the correct value. As apparent from Figure 4, the bias of the 1000 draws using residual bootstrap was larger than that of our Bayesian approach, and the attempted bias correction was not able to compensate for it fully.

Next, we investigated the posterior distribution of the crossing angle when using Constrained Spherical Deconvolution (CSD) (Tournier et al., 2007). The crossing angle was

estimated as the angle between the two largest peaks on the fiber orientation distribution function, provided this angle was at least 25° . We found the peak detection to be more accurate when using only single-shell data. Consequently, we used only the $b = 3000 \text{ s/mm}^2$ shell to fit a CSD model of order 10 to each simulated measurement. We set the regularization strength λ to 5, and the amplitude threshold τ (below which the corresponding fiber orientation density is assumed to be zero) to 0.1. The single tensor model was used as the response function. With a crossing angle of 45° and 60° , the peak detection algorithm successfully detected two distinct peaks in 97.4% and 100% out of the 1000 cases, respectively. The average estimate of the crossing angle was 47.2° and 60.0° , respectively. Figure 5 shows the resulting P-P plots, and compares with 1000 draws using residual bootstrap.

3.2. In vivo data

We show two types of results using in vivo data: first, maps of FA and RTOP with corresponding uncertainties for two healthy adults from the the Human Connectome Project (HCP) (Van Essen et al., 2013); second, the application of our Bayesian approach to a realistic group analysis scenario where we perform voxelwise comparisons of FA between controls and individuals diagnosed with schizophrenia.

3.2.1. Maps of quantitative features and corresponding uncertainties—We used the freely available MGH adult diffusion dataset from the Human Connectome Project^{2,3} (HCP) (Van Essen et al., 2013). The subjects are healthy adults, scanned on a customized Siemens 3T Connectome scanner (Keil et al., 2013; Setsompop et al., 2013) using a Stejskal-Tanner type diffusion weighted spin-echo sequence. The downloaded data have already been corrected for gradient nonlinearities, subject motion and eddy currents (Andersson and Sotiropoulos, 2016). Diffusion measurements were at four b-value shells: 1000, 3000, 5000, 10000 s/mm^2 . The corresponding number of gradient orientations were 64, 64, 128 and 256. A total of 40 non-diffusion weighted ($b = 0$) images were acquired at regular intervals.

We analyzed subjects mgh_1007 and mgh_1010. First, we used the same procedure as in the single tensor simulation to fit a diffusion tensor model (weighted least-squares on $b = 1000 \text{ s/mm}^2$). We estimated the posterior distribution of fractional anisotropy in each voxel from a draw of 1000 samples. The mean and the interquartile range of the fractional anisotropy is shown in Figure 6. Second, we used all measurements with $b = 3000 \text{ s/mm}^2$ to fit an anisotropic MAP-MRI model of order 6 with Laplacian regularization determined by generalized cross validation. Figure 7 shows the mean and interquartile range of the return to origin probability in each voxel. On a standard laptop, the computation time needed to estimate FA by fitting a DTI model to each voxel in a slice, such as in Figure 6, was about 3 s whereas sampling in order to estimate the interquartile range required about 90 s. Similarly, estimating RTOP by fitting a MAP-MRI model required about 220 s whereas

²<http://www.humanconnectome.org/documentation/MGH-diffusion/>

³<https://ida.loni.usc.edu/login.jsp>

obtaining the corresponding, closed-form, estimates of the interquartile range only required 3 s.

3.2.2. Group analysis of FA differences in schizophrenia—To apply our Bayesian approach to a realistic group analysis scenario, we downloaded diffusion data from the OpenfMRI project (Poldrack et al., 2013). Specifically, we downloaded data from the Consortium for Neuropsychiatric Phenomics (CNP) (Poldrack et al., 2016)⁴, which provides data from 138 healthy controls, 58 individuals diagnosed with schizophrenia, 49 individuals with bipolar disorder and 45 individuals with ADHD. Out of these subjects, we selected 20 healthy controls and 20 individuals with schizophrenia. The CNP dataset contains anatomical T_1 volumes, resting state fMRI, task fMRI as well as diffusion MRI. The diffusion data contain 64 diffusion weighted volumes ($b = 1000 \text{ s/mm}^2$) and 1 volume without diffusion weighting ($b = 0$), and the voxels have a size of $2 \times 2 \times 2 \text{ mm}^3$.

We corrected the diffusion data for head motion and eddy currents using the function `eddy_correct` in FSL (as it is much faster compared to the more advanced `eddy` function (Andersson and Sotiropoulos, 2016)). Our Bayesian approach was then used to generate a posterior distribution of FA for each subject and each voxel inside the brain. The standard TBSS (Smith et al., 2006) pipeline was used to generate a mean FA skeleton, and 1000 draws of FA for each of the 40 subjects were then projected onto this skeleton (by modifying the script `tbss_non_FA`). Two Bayesian group analyses were performed to investigate if there are differences in FA between the two groups. In the first group analysis each subject was treated equally, and each sample in the group posterior $p(\overline{FA}^{\text{controls}} - \overline{FA}^{\text{schizophrenics}} | y)$ was calculated using voxelwise, unweighted, means for each group

$$\overline{FA}_v^{\text{group}} = \frac{1}{N^{\text{group}}} \sum_{i \in \text{group}} FA_{i,v}, \quad (27)$$

where $FA_{i,v}$ represents the posterior distribution for subject i and voxel v , and N^{group} is the number of subjects in that group. In the second group analysis, we used the uncertainty quantification provided by our Bayesian approach to weight each subject and voxel. We defined the weight $w_{i,v}$ as the reciprocal of the (sample) standard deviation of the corresponding FA posterior distribution (Beckmann et al., 2003), i.e. as

$$w_{i,v} = \frac{1}{\text{SD}(FA_{i,v})}. \quad (28)$$

A weighted mean was calculated for each group, according to

⁴<https://openfMRI.org/dataset/ds000030/>

$$\widehat{FA}_v^{\text{group}} = \frac{\sum_{i \in \text{group}} FA_{i,v} \cdot w_{i,v}}{\sum_{i \in \text{group}} w_{i,v}}. \quad (29)$$

Then, the differences between the weighted group means were considered as samples from the group posterior.

Finally, a Bayesian t-score (group posterior mean divided by group posterior standard deviation) was calculated in each voxel, since posterior probability maps are rather sensitive to the effect size threshold. For example, calculating the posterior probability $p((\widehat{FA}^{\text{controls}} - \widehat{FA}^{\text{schizophrenics}}) > 0|y)$ for every voxel often leads to a probability close to 1 in a very large proportion of the voxels, which is not informative. An effect size threshold α can be used to provide more informative posterior probability maps, by instead calculating $p((\widehat{FA}^{\text{controls}} - \widehat{FA}^{\text{schizophrenics}}) > \alpha|y)$, but then the question is what threshold to use. We therefore show maps of Bayesian t-scores, instead of posterior probability maps.

Figure 8 shows the mean weights for the healthy controls and the individuals diagnosed with schizophrenia, as well as their difference. Figure 9 shows results from the unweighted (regular) as well as the weighted group analysis. Figure 10 shows a (hopefully) pedagogical comparison between the unweighted and the weighted group analysis, for one voxel where the difference between the two approaches is rather large. The weighted group analysis downweights the influence of a subject that has a much higher uncertainty than the others.

4. Discussion

4.1. Interpretation of the results

For the single tensor simulation fitted with DTI, the P-P plot in Figure 1, for the mean diffusivity, shows excellent agreement between theory and observation. This is expected due to the almost perfect match between model and (simulated) measurements — the signal is generated by a DTI model and for $b = 1000$ s/mm² the signal-to-noise ratio is roughly between 5 and 15, meaning that the likelihood (i.e. measurement distribution) is well approximated as Gaussian (Gudbjartsson and Patz, 1995; Salvador et al., 2005). The P-P plot in Figure 2, for the fractional anisotropy, reflects the previously established fact that there is an increasing upward bias of the estimated FA as the true FA decreases (Bastin et al., 1998; Farrell et al., 2007).

From Figures 1 and 2 we observe that, in the single tensor simulations, the results of our Bayesian approach and those using residual bootstrap are virtually indistinguishable. The extensive computational requirements of bootstrap methods is one of their main drawbacks. It is thus encouraging that our Bayesian approach offers substantial speed-ups compared to residual bootstrap using 1000 draws—around 20 times in the FA case (that requires sampling) and 200 times in the MD case (using the closed-form expression).

The issue with bias is also present in the RTOP estimation, as can be seen in Figure 4. After subtracting the bias from the theoretical quantiles, we can see, in Figures 4b and 4d, that

there is good agreement between the bias-corrected theoretical distributions and the observed ones. This indicates that the shapes of the posterior distributions are similar to the correct ones, but that their means differ. In other words, the Bayesian uncertainty quantification is accurate although the point estimate (the mean) is not. The results using residual bootstrap, however, exhibit a larger bias which was not fully compensated by the attempted bias correction. We hypothesize that this is because the signal-to-noise ratio differs in the different shells, so that the residuals are not interchangeable (as assumed in residual bootstrap). A potential solution would thus be to only resample within the same shell. There are also more general approaches that could reduce the bias, for instance to denoise the data (Wiest-Daesslé et al., 2008) or to transform the signal distribution from Rician to Gaussian (Koay et al., 2009).

In the crossing angle estimation using CSD, Figure 5, we see that the results of our approach and residual bootstrap are different but of comparable, fairly high, quality.

For the scalar indices extracted from in vivo data and shown in Figures 6 and 7, we find that the noise levels that can be appreciated qualitatively from the mean images correspond well with larger interquartile ranges. We note that the FA uncertainty appears to increase with the distance from the head's surface. For the RTOP uncertainty, on the other hand, the predominant effect seems to be that the uncertainty correlates with the value of the RTOP. In other words, the RTOP estimates appear to be heteroscedastic.

Although we have not made any effort to optimize our implementation for performance, we argue that the additional computational cost of uncertainty quantification is within reason. In particular, if the model is complex and the quantity of interest is an affine function of the coefficients — as was the case with the RTOP estimation using MAP-MRI — then the computational overhead of the uncertainty quantification is almost negligible. The FA estimation using DTI, on the other hand, is an example where fitting the model is cheap but estimating the uncertainty requires sampling.

4.2. Group analysis

As can be seen in Figure 8, the subject and voxel specific weights differ between the two subject groups. The controls in general have a higher weight (i.e. lower uncertainty). This can for example be explained by differences in the amount of head motion (Yendiki et al., 2014), or that different MR scanners have different noise levels (at least two MR scanners were used in the CNP study). Existing approaches for group analyses in diffusion MRI, such as TBSS (Smith et al., 2006), do not use the uncertainty, and can for example not downweight subjects with a high degree of head motion. It is possible that using the more advanced eddy function (Andersson and Sotiropoulos, 2016) for preprocessing would lead to more similar weights for all subjects.

Figure 9 indicates minor differences between the unweighted (regular) and the weighted group analysis. The differences are mostly located in the outermost parts of the skeleton, which is consistent with the explanation that they are induced by artefacts. However, as illustrated by Figure 10, a high uncertainty for a single subject can lead to a rather big difference between the two approaches. Both the unweighted and the weighted group

analyses indicate a mix of stronger and weaker FA between the groups, while previous schizophrenia studies mainly have found stronger FA in healthy controls (Kubicki et al., 2007). The standard TBSS group analysis (based on a permutation test) did not reveal any significant group differences, which for example can be explained by low statistical power. The aim of our study was not group differences per se, but to compare the unweighted and the weighted approach.

A possible explanation why weighted group analyses are uncommon in the diffusion MRI field, compared to the fMRI field (Chen et al., 2012; Woolrich et al., 2004), is the higher computational complexity. To generate 1000 draws of FA for one subject took approximately 30 minutes, giving a total of 20 hours for 40 subjects (which is fast compared to Bayesian methods using Markov Chain Monte Carlo (MCMC), often requiring many hours for a single subject (Wegmann et al., 2017)). The TBSS pipeline projects the FA values onto a skeleton where the voxels are $1 \times 1 \times 1$ mm. In our case this resulted in a total of 40 000 volumes with $182 \times 218 \times 182$ voxels each, requiring about 1200 GB of storage. In fMRI the standard is to register each brain to a brain template with an isotropic resolution of 2 or 3 mm, which significantly reduces the memory consumption. Furthermore, it is in fMRI sufficient to store the mean and variance for each voxel (if a normal distribution is assumed), while diffusion MRI may require storing many draws from a more complicated distribution (e.g. FA). The FA distribution in each voxel can be fitted to a beta distribution, as done in Figure 10, which makes it sufficient to store only the two shape parameters in every voxel. Similarly to Bayesian methods for tractography, in FSL based on MCMC (Behrens et al., 2007), graphics processing units (GPUs) can be used to speedup subject level as well as group level analyses (Eklund et al., 2013; Harms et al., 2017; Hernandes et al., 2013).

4.3. General reflections

The class of models fitted with linear least-squares encompasses many of the most popular models. Nevertheless, it would be possible to conduct the same type of probabilistic reasoning for many of the outstanding models, e.g. those using l_1 -penalization (which would correspond to having a Laplace prior on the coefficients). The catch, however, is that it would no longer be possible to arrive at a closed form expression for the posterior distribution. However, it would still be possible to estimate the posterior by Monte Carlo sampling or approximate methods such as variational Bayes (Gelman et al., 2013). It could then be viewed as part of a general framework called probabilistic programming (Ghahramani, 2015).

In addition to using uncertainty in group analyses and tractography, as we have described before, there are other applications where it could prove useful. One example is in the calculation of the ensemble averaged propagator (EAP), e.g. as proposed in (Sjölund et al., 2017).

5. Conclusions

We make a Bayesian reinterpretation of some of the most popular algorithms for signal estimation in dMRI. In this interpretation, the commonly used point estimate only constitute

the mean of a multivariate probability distribution. With access to this distribution it is, in principle, possible to quantify the uncertainty of any derived quantity. In particular, for quantities that are affine functions of the coefficients the posterior can be expressed in closed-form. As an example of a dMRI application that we think would benefit in particular from this type of uncertainty quantification, we show results from a realistic group analysis scenario, indicating that a weighted group analysis can indeed downweight subjects with a higher uncertainty.

Acknowledgements

This study was supported by the Swedish Foundation for Strategic Research (grant AM13–0090), the Swedish Research Council CADICS Linneaus research environment, the Swedish Research Council (grants 2012–4281, 2013–5229, 2015–05356 and 2016–04482), Linköping University Center for Industrial Information Technology (CENIT), VINNOVA/ITEA3 BENEFIT (grant 2014–00593), and National Institutes of Health (grants P41EB015902, R01MH074794, P41EB015898), and the Knut and Alice Wallenberg Foundation project “Seeing Organ Function”.

Data collection and sharing for this project was provided by the Human Connectome Project (HCP; Principal Investigators: Bruce Rosen, M.D., Ph.D., Arthur W. Toga, Ph.D., Van J. Weeden, MD). HCP funding was provided by the National Institute of Dental and Craniofacial Research (NIDCR), the National Institute of Mental Health (NIMH), and the National Institute of Neurological Disorders and Stroke (NINDS). HCP data are disseminated by the Laboratory of Neuro Imaging at the University of Southern California.

Appendix A. Marginalization over the residual variance

In this section we first show that by placing an inverse-Gamma prior over the residual variance σ^2 , the posterior becomes a multivariate t -distribution. Then we show that it is possible to choose the prior hyperparameters such that the covariance and the degrees of freedom remain the same as before putting a prior over σ^2 , ultimately leading to the expressions in equations (18) and (19).

The inverse-Gamma distribution is a convenient choice because it is the conjugate prior for σ^2 . Its probability distribution is

$$p(\sigma^2) = IG(\alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \frac{1}{(\sigma^2)^{\alpha+1}} \exp\left(-\frac{\beta}{\sigma^2}\right), \quad (\text{A.1})$$

where $\Gamma(\cdot)$ is the Gamma function.

The joint posterior of $\mathbf{c}$ and σ^2 can be computed by combining the Gaussian likelihood (5), the Gaussian prior (6) on the coefficients and the Inverse-Gamma prior (A.1) on σ^2 ,

$$\begin{aligned}
p(\mathbf{c}, \sigma^2 | \mathbf{y}, \mathbf{x}) &\propto p(\mathbf{y} | \mathbf{c}, \sigma^2, \mathbf{x}) p(\mathbf{c} | \sigma^2) p(\sigma^2) \\
&= \frac{|W|^{1/2}}{(2\pi\sigma^2)^{n/2}} \exp\left(-\frac{1}{2\sigma^2} (\mathbf{y} - \Phi \mathbf{c})^T W (\mathbf{y} - \Phi \mathbf{c})\right) \\
&\cdot \frac{|A|^{1/2}}{(2\pi\sigma)^{d/2}} \exp\left(-\frac{1}{2\sigma^2} \mathbf{c}^T A \mathbf{c}\right) \cdot \frac{\beta^\alpha}{\Gamma(\alpha)} \frac{1}{(\sigma^2)^{\alpha+1}} \exp\left(-\frac{\beta}{\sigma^2}\right) \\
&\propto (\sigma^2)^{-\left(\alpha+1+\frac{n}{2}+\frac{d}{2}\right)} \exp\left(-\frac{1}{\sigma^2} \left(\beta + \frac{1}{2} \mathbf{y}^T W \mathbf{y} + \frac{1}{2} (\mathbf{c}^T Q \mathbf{c} - 2 \mathbf{c}^T \Phi^T W \mathbf{y})\right)\right) \\
&= (\sigma^2)^{-\left(\alpha+1+\frac{n}{2}+\frac{d}{2}\right)} \exp\left(-\frac{1}{\sigma^2} \left(\beta + \frac{1}{2} (\mathbf{y}^T W \mathbf{y} - \mathbf{y}^T W \Phi Q^{-1} \Phi^T W \mathbf{y})\right)\right) \\
&\cdot \exp\left(-\frac{1}{2\sigma^2} (\mathbf{c} - Q^{-1} \Phi^T W \mathbf{y})^T Q (\mathbf{c} - Q^{-1} \Phi^T W \mathbf{y})\right) \\
&= \mathcal{N}(\mathbf{c} | \boldsymbol{\mu}, \sigma^2 Q^{-1}) IG(\sigma^2 | \alpha_*, \beta_*),
\end{aligned}$$

Where

$$\alpha_* = \alpha + \frac{n}{2}, \quad (\text{A.2})$$

$$\beta_* = \beta + \frac{1}{2} \mathbf{y}^T W (\mathbf{y} - \Phi \boldsymbol{\mu}). \quad (\text{A.3})$$

Now, we obtain the marginal posterior distribution by integrating over σ^2 ,

$$p(\mathbf{c} | \mathbf{y}, \mathbf{x}) = \int p(\mathbf{c}, \sigma^2 | \mathbf{y}, \mathbf{x}) d(\sigma^2) \quad (\text{A.4})$$

$$= \int \mathcal{N}(\mathbf{c} | \boldsymbol{\mu}, \Sigma) IG(\sigma^2 | \alpha_*, \beta_*) d(\sigma^2) \quad (\text{A.5})$$

$$= \frac{\Gamma(\alpha_* + \frac{d}{2})}{\Gamma(\alpha_*)} \frac{|Q|^{1/2}}{(2\pi\beta_*)^{d/2}} \left(1 + \frac{(\mathbf{c} - \boldsymbol{\mu})^T Q (\mathbf{c} - \boldsymbol{\mu})}{2\beta_*}\right)^{-\left(\alpha_* + \frac{d}{2}\right)} \quad (\text{A.6})$$

which is in fact a multivariate t -distribution (Kotz and Nadarajah, 2004; Roth, 2012) with degrees of freedom $\nu = 2\alpha_*$, mean vector $\boldsymbol{\mu}$ and correlation matrix $R = (\beta_*/\alpha_*) Q^{-1}$. For this result to be usable we must specify the prior hyper-parameters α and β . In the spirit of what

is commonly referred to as empirical Bayes, we determine these hyperparameters by requiring that the covariance and the degrees of freedom in equation (18) remain the same as before putting a distribution over σ^2 . Setting $\nu = 2\alpha_*$, we may write the covariance in (18) is

$$\frac{\nu}{\nu-2}R = \frac{\beta_*}{\alpha_*-1}Q^{-1} = E[\sigma^2]Q^{-1}. \quad (\text{A.7})$$

To get correspondence with the covariance in equation (8), we require that $E[\sigma^2] = \hat{\sigma}^2$. Combining this with equations (A.3) gives

$$\beta = -\frac{1}{2}\left(\frac{1}{\alpha_*}\mathbf{y} + \frac{\alpha_*-1}{\alpha_*}\hat{\mathbf{y}}\right)^T W(\mathbf{y} - \hat{\mathbf{y}}). \quad (\text{A.8})$$

This means that we can express the degrees of freedom and correlation matrix of the multivariate t -distribution in terms of ν and $\hat{\sigma}^2$ as

$$\begin{aligned} \nu &= \|I - H\|_{\text{Fro}}^2, \\ R &= \frac{\nu-2}{\nu}\hat{\sigma}^2Q^{-1}. \end{aligned} \quad (\text{A.9})$$

References

- Andersson JL, Sotiropoulos SN, 2015 Non-parametric representation and prediction of single-and multi-shell diffusion-weighted MRI data using Gaussian processes. *NeuroImage* 122, 166–176. [PubMed: 26236030]
- Andersson JL, Sotiropoulos SN, 2016 An integrated approach to correction for off-resonance effects and subject movement in diffusion MR imaging. *Neuroimage* 125, 1063–1078. [PubMed: 26481672]
- Asselmlal H-E, Tschumperlé D, Brun L, 2009 Efficient and robust computation of PDF features from diffusion MR signal. *Medical image analysis* 13 (5), 715–729. [PubMed: 19665917]
- Basser PJ, Mattiello J, LeBihan D, 1994 MR diffusion tensor spectroscopy and imaging. *Biophysical journal* 66 (1), 259. [PubMed: 8130344]
- Bastin ME, Armitage PA, Marshall I, 1998 A theoretical study of the effect of experimental noise on the measurement of anisotropy in diffusion imaging. *Magnetic resonance imaging* 16 (7), 773–785. [PubMed: 9811143]
- Beckmann CF, Jenkinson M, Smith SM, 2003 General multilevel linear modeling for group analysis in FMRI. *NeuroImage* 20 (2), 1052–063. [PubMed: 14568475]
- Behrens TE, Berg HJ, Jbabdi S, Rushworth MF, Woolrich MW, 2007 Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? *Neuroimage* 34 (1), 144–155. [PubMed: 17070705]
- Behrens TE, Woolrich MW, Jenkinson M, Johansen-Berg H, Nunes RG, Clare S, Matthews PM, Brady JM, Smith SM, 2003 Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magnetic resonance in medicine* 50 (5), 1077–1088. [PubMed: 14587019]
- Berman JI, Chung S, Mukherjee P, Hess CP, Han ET, Henry RG, 2008 Probabilistic streamline q-ball tractography using the residual bootstrap. *Neuroimage* 39 (1), 215–222. [PubMed: 17911030]
- Bishop CM, 2006 Pattern recognition and machine learning Springer.

- Caruyer E, Deriche R, 2012 Diffusion MRI signal reconstruction with continuity constraint and optimal regularization. *Medical image analysis* 16 (6), 1113–1120. [PubMed: 22854036]
- Chen G, Saad ZS, Nath AR, Beauchamp MS, Cox RW, 2012 FMRI group analysis combining effect estimates and their variances. *Neuroimage* 60 (1), 747–765. [PubMed: 22245637]
- Chung S, Lu Y, Henry RG, 2006 Comparison of bootstrap approaches for estimation of uncertainties of DTI parameters. *NeuroImage* 33 (2), 531–541. [PubMed: 16938472]
- Davison AC, Hinkley DV, 1997 *Bootstrap methods and their application* Vol. 1 Cambridge university press.
- Descoteaux M, Angelino E, Fitzgibbons S, Deriche R, 2006 Apparent diffusion coefficients from high angular resolution diffusion imaging: Estimation and applications. *Magnetic Resonance in Medicine* 56 (2), 395–410. [PubMed: 16802316]
- Descoteaux M, Angelino E, Fitzgibbons S, Deriche R, 2007 Regularized, fast, and robust analytical Q-ball imaging. *Magnetic resonance in medicine* 58 (3), 497–510. [PubMed: 17763358]
- Eklund A, Dufort P, Forsberg D, LaConte S, 2013 Medical image processing on the GPU - Past, present and future. *Medical Image Analysis* 17, 1073–1094. [PubMed: 23906631]
- Farrell JA, Landman BA, Jones CK, Smith SA, Prince JL, van Zijl P, Mori S, 2007 Effects of signal-to-noise ratio on the accuracy and reproducibility of diffusion tensor imaging-derived fractional anisotropy, mean diffusivity, and principal eigenvector measurements at 1.5 T. *Journal of Magnetic Resonance Imaging* 26 (3), 756–767. [PubMed: 17729339]
- Fick RH, Wassermann D, Caruyer E, Deriche R, 2016 MAPL: Tissue microstructure estimation using Laplacian-regularized MAP-MRI and its application to HCP data. *NeuroImage* 134, 365–385. [PubMed: 27043358]
- Garyfallidis E, Brett M, Amirbekian B, Rokem A, Van Der Walt S, Descoteaux M, Nimmo-Smith I, Contributors D, 2014 Dipy, a library for the analysis of diffusion MRI data. *Frontiers in neuroinformatics* 8.
- Gelman A, Carlin JB, Stern HS, Dunson DB, Vehtari A, Rubin DB, 2013 *Bayesian Data Analysis* CRC Press.
- Ghahramani Z, 2015 Probabilistic machine learning and artificial intelligence. *Nature* 521 (7553), 452. [PubMed: 26017444]
- Gu X, Sidén P, Wegmann B, Eklund A, Villani M, Knutsson H, 2017 Bayesian diffusion tensor estimation with spatial priors. In: *International Conference on Computer Analysis of Images and Patterns* Springer, Cham, pp. 372–383.
- Gudbjartsson H, Patz S, 1995 The Rician distribution of noisy MRI data. *Magnetic resonance in medicine* 34 (6), 910–914. [PubMed: 8598820]
- Harms R, Fritz F, Tobisch A, Goebel R, Roebroek A, 2017 Robust and fast nonlinear optimization of diffusion MRI microstructure models. *NeuroImage* 155, 82–96. [PubMed: 28457975]
- Haroon HA, Morris DM, Embleton KV, Alexander DC, Parker GJ, 2009 Using the model-based residual bootstrap to quantify uncertainty in fiber orientations from q-ball analysis. *IEEE transactions on medical imaging* 28 (4), 535–550. [PubMed: 19272997]
- Hastie T, Tibshirani R, Friedman J, 2001 *The elements of statistical learning* Vol. 1 Springer series in statistics New York.
- Hernandes M, Guerrero G, Cecilia J, Garcia J, Inuggi A, Jbabdi S, Behrens T, Sotiropoulos SN, 2013 Accelerating fibre orientation estimation from diffusion weighted magnetic resonance imaging using GPUs. *PLOS ONE* 8 (4), 1–13.
- Hess CP, Mukherjee P, Han ET, Xu D, Vigneron DB, 2006 Q-ball reconstruction of multimodal fiber orientations using the spherical harmonic basis. *Magnetic Resonance in Medicine* 56 (1), 104–117. [PubMed: 16755539]
- Jensen JH, Helpert JA, Ramani A, Lu H, Kaczynski K, 2005 Diffusional kurtosis imaging: The quantification of non-Gaussian water diffusion by means of magnetic resonance imaging. *Magnetic resonance in medicine* 53 (6), 1432–1440. [PubMed: 15906300]
- Jeurissen B, Leemans A, Jones DK, Tournier J-D, Sijbers J, 2011 Probabilistic fiber tracking using the residual bootstrap with constrained spherical deconvolution. *Human brain mapping* 32 (3), 461–479. [PubMed: 21319270]

- Johansen-Berg H, Rushworth MF, 2009 Using diffusion imaging to study human connective anatomy. *Annual review of neuroscience* 32, 75–94.
- Jones DK, 2003 Determining and visualizing uncertainty in estimates of fiber orientation from diffusion tensor MRI. *Magnetic Resonance in Medicine* 49 (1), 7–12. [PubMed: 12509814]
- Jones DK, 2008 Tractography gone wild: probabilistic fibre tracking using the wild bootstrap with diffusion tensor MRI. *IEEE transactions on medical imaging* 27 (9), 1268–1274. [PubMed: 18779066]
- Keil B, Blau JN, Biber S, Hoecht P, Tountcheva V, Setsompop K, Triantafyllou C, Wald LL, 2013 A 64-channel 3T array coil for accelerated brain MRI. *Magnetic resonance in medicine* 70 (1), 248–258. [PubMed: 22851312]
- Koay CG, Chang L-C, Pierpaoli C, Basser PJ, 2007 Error propagation framework for diffusion tensor imaging via diffusion tensor representations. *IEEE transactions on medical imaging* 26 (8), 1017–1034. [PubMed: 17695123]
- Koay CG, Özarslan E, Basser PJ, 2009 A signal transformational framework for breaking the noise floor and its applications in MRI. *Journal of magnetic resonance* 197 (2), 108–119. [PubMed: 19138540]
- Kotz S, Nadarajah S, 2004 *Multivariate t-distributions and their applications* Cambridge University Press.
- Kubicki M, McCarley R, Westin C-F, Park H-J, Maier S, Kikinis R, Jolesz FA, Shenton ME, 2007 A review of diffusion tensor imaging studies in schizophrenia. *Journal of psychiatric research* 41 (1), 15–30. [PubMed: 16023676]
- Loader C, 1999 *Local regression and likelihood* Springer-Verlag New York.
- Özarslan E, Koay C, Shepherd T, Komlos M, rfano lu M, Pierpaoli C, Basser P, 2013 Mean apparent propagator (MAP) MRI: a novel diffusion imaging method for mapping tissue microstructure. *NeuroImage* 78, 16–32. [PubMed: 23587694]
- Özarslan E, Vemuri BC, Mareci TH, 2005 Generalized scalar measures for diffusion MRI using trace, variance, and entropy. *Magnetic resonance in Medicine* 53 (4), 866–876. [PubMed: 15799039]
- Pierpaoli C, Basser PJ, 1996 Toward a quantitative assessment of diffusion anisotropy. *Magnetic resonance in Medicine* 36 (6), 893–906. [PubMed: 8946355]
- Poldrack R, Barch D, Mitchell J, Wager T, Wagner A, Devlin J, Cumba C, Koyejo O, Milham M, 2013 Toward open sharing of task-based fMRI data: the OpenfMRI project. *Frontiers in Neuroinformatics* 7 (12).
- Poldrack R, Congdon E, Triplett W, Gorgolewski K, Karlsgodt K, Mum-ford J, Sabb F, Freimer N, London E, Cannon T, Bilder R, 2016 A phenome-wide examination of neural and cognitive function. *Scientific Data* 3 (160110).
- Rawlings JO, Pantula SG, Dickey DA, 2001 *Applied regression analysis: a research tool* Springer Science & Business Media.
- Roth M, 2012 On the multivariate t distribution. Tech. rep, Linköping University, Department of Electrical Engineering.
- Salvador R, Peña A, Menon DK, Carpenter TA, Pickard JD, Bull-more ET, 2005 Formal characterization and extension of the linearized diffusion tensor model. *Human brain mapping* 24 (2), 144–155. [PubMed: 15468122]
- Setsompop K, et al., 2013 Pushing the limits of in vivo diffusion MRI for the Human Connectome Project. *Neuroimage* 80, 220–233. [PubMed: 23707579]
- Sjölund J, Eklund A, Özarslan E, Knutsson H, 2017 Gaussian process regression can turn non-uniform and undersampled diffusion MRI data into diffusion spectrum imaging. In: *IEEE 14th International Symposium on Biomedical Imaging (ISBI) IEEE*, pp. 778–782.
- Smith SM, Jenkinson M, Johansen-Berg H, Rueckert D, Nichols TE, Mackay CE, Watkins KE, Ciccarelli O, Cader MZ, Matthews PM, et al., 2006 Tract-based spatial statistics: voxelwise analysis of multi-subject diffusion data. *Neuroimage* 31 (4), 1487–1505. [PubMed: 16624579]
- Stejskal E, Tanner J, 1965 Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient. *The journal of chemical physics* 42 (1), 288–292.

- Tournier J-D, Calamante F, Connelly A, 2007 Robust determination of the fibre orientation distribution in diffusion MRI: non-negativity constrained super-resolved spherical deconvolution. *Neuroimage* 35 (4), 1459–1472. [PubMed: 17379540]
- Tuch DS, 2004 Q-ball imaging. *Magnetic resonance in medicine* 52 (6), 1358–1372. [PubMed: 15562495]
- Van Essen D, Smith S, Barch D, Behrens T, Yacoub E, Ugurbil K, 2013 The WU-Minn Human Connectome Project: an overview. *Neuroimage* 80, 62–79. [PubMed: 23684880]
- Veraart J, Sijbers J, Sunaert S, Leemans A, Jeurissen B, 2013 Weighted linear least squares estimation of diffusion MRI parameters: strengths, limitations, and pitfalls. *NeuroImage* 81, 335–346. [PubMed: 23684865]
- Wegmann B, Eklund A, Villani M, 2017 Bayesian Rician regression for neuroimaging. *Frontiers in Neuroscience* 11, 586. [PubMed: 29104529]
- Westin CF, Knutsson H, Pasternak O, Szczepankiewicz F, Özarslan E, van Westen D, Mattisson C, Bogren M, O'Donnell LJ, Kubicki M, et al., 2016 Q-space trajectory imaging for multidimensional diffusion MRI of the human brain. *Neuroimage* 135, 345–362. [PubMed: 26923372]
- Whitcher B, Tuch DS, Wisco JJ, Sorensen AG, Wang L, 2008 Using the wild bootstrap to quantify uncertainty in diffusion tensor imaging. *Human brain mapping* 29 (3), 346–362. [PubMed: 17455199]
- Wiest-Daesslé N, Prima S, Coupé P, Morrissey SP, Barillot C, 2008 Rician noise removal by non-local means filtering for low signal-to-noise ratio MRI: applications to DT-MRI. In: *International Conference on Medical Image Computing and Computer-assisted Intervention* Springer, pp. 171–179.
- Wilk MB, Gnanadesikan R, 1968 Probability plotting methods for the analysis of data. *Biometrika* 55 (1), 1–17. [PubMed: 5661047]
- Woolrich MW, Behrens TE, Beckmann CF, Jenkinson M, Smith SM, 2004 Multilevel linear modelling for fMRI group analysis using Bayesian inference. *Neuroimage* 21 (4), 1732–1747. [PubMed: 15050594]
- Yendiki A, Koldewyn K, Kakunoori S, Kanwisher N, Fischl B, 2014 Spurious group differences due to head motion in a diffusion MRI study. *NeuroImage* 88, 79–90. [PubMed: 24269273]
- Yuan Y, Zhu H, Ibrahim JG, Lin W, Peterson BS, 2008 A note on the validity of statistical bootstrapping for estimating the uncertainty of tensor parameters in diffusion tensor images. *IEEE transactions on medical imaging* 27 (10), 1506–1514. [PubMed: 18815102]

Figure 1:

P-P plots of mean diffusivity (MD) for 1000 simulated measurements from three single tensor models with different fractional anisotropy (FA) values. Mean diffusivity was estimated using a DTI model fitted with weighted least-squares. Averaging over the simulations, the P-P plots show how often the correct value is smaller than the quantiles corresponding to the theoretically expected probabilities for that property given a simulated measurement.

Figure 2:

P-P plots of fractional anisotropy (FA) for 1000 simulated measurements from three single tensor models with different FA values. FA was estimated using a DTI model fitted with weighted least-squares. Averaging over the simulations, the P-P plots show how often the correct value is smaller than the quantiles corresponding to the theoretically expected probabilities for that property given a simulated measurement.

Figure 3:

Average ratio of estimated to true values of the return to origin probability (RTOP) in 1000 simulated measurements following Gaussian or Rician distributions with scale $\sigma = 0.05S_0$. The signal was generated from a 45° double tensor model and was fitted using anisotropic MAPL models of orders 4, 6 and 8 with regularization strength determined by generalized cross-validation. The experimental parameters were the same as in the Human Connectome Project, except that the maximum b -value, b_{max} , used in the fitting was gradually increased.

Figure 4:

P-P plots of return to origin probability for 1000 simulated measurements from 45° and 60° double tensor models fitted with anisotropic MAPL models of order 4. Averaging over the simulations, the P-P plots show how often the correct value is smaller than the quantiles corresponding to the theoretically expected probabilities for that property given a simulated measurement. In (b) and (d), the bias (average error of the mean) was subtracted from the theoretical quantiles before comparing with the correct value; this corresponds to shifting the mean of the posterior distribution to the correct value.

Figure 5:

P-P plots of the crossing angle for 1000 simulated measurements from double tensor models fitted with constrained spherical deconvolution (CSD) models of order 10. Averaging over the simulations, the P-P plots show how often the correct value is smaller than the quantiles corresponding to the theoretically expected probabilities for that property given a simulated measurement.

Figure 6: Mean and interquartile range of the fractional anisotropy (FA) estimated using a DTI model fitted with weighted least-squares. The figures correspond to subjects mgh_1007 (top) and mgh_1010 (bottom) in the MGH adult diffusion dataset, but using only measurements with b 1000 s/mm²

Figure 7: Mean and interquartile range of the return to origin probability (RTOP) estimated using an anisotropic MAP-MRI model of order 6 with Laplacian regularization. The figures correspond to subjects mgh_1007 (top) and mgh_1010 (bottom) in the MGH adult diffusion dataset, but using only measurements with $b \leq 3000 \text{ s/mm}^2$.

Figure 8:

A comparison of weights (calculated as the reciprocal of the posterior standard deviation in each voxel) for healthy controls and individuals diagnosed with schizophrenia. Figures a) and b) show the mean weights over the 20 subjects in each group, while figure c) shows the difference between the mean weights.

Figure 9:

Results from the two Bayesian group analyses; a) unweighted (regular), b) weighted and c) weighted minus unweighted. All maps have arbitrarily been thresholded at a Bayesian t-score of 2.5. The color coding refers to the Bayesian t-score. By downweighting subjects with a high uncertainty, the weighted approach can produce different group results, especially for voxels that are more sensitive to different forms of artefacts (e.g. voxels close to the edge of the brain) and for voxels located in smaller fiber bundles (where FA may be more uncertain). This can be seen in c), as the major differences between the two approaches are located in the outer parts of the skeleton.

(a) FA posterior distributions for all 40 subjects in one voxel (a beta distribution was fitted to the posterior of each subject, using the 1000 draws). One subject with schizophrenia has a much higher uncertainty compared to all other subjects, but this uncertainty is ignored in the unweighted group analysis.

(b) Posterior distributions of the group difference in FA (controls - schizophrenia) in one voxel, for the unweighted and the weighted group analysis (a normal distribution was fitted to each posterior, using the 1000 draws).

Figure 10:

A comparison of the unweighted and the weighted group analysis in a voxel where the difference between the two approaches is rather large. Subjects with a higher uncertainty will be downweighted in the weighted group analysis, which can lead to a rather different posterior of the group difference.

Table 1:

An assortment of linear models used in dMRI. We have used the notation $\mathbf{q} = q\hat{\mathbf{q}}$ where applicable. For clarity, some details have been omitted.

Method	y	$\phi(\mathbf{x})$	$\mathbf{c}$	$\mathbf{W}$	Λ
Diffusion Tensor Imaging (DTI) (weighted least-squares) (Basser et al., 1994; Salvador et al., 2005)	$\log S(\mathbf{q})$	$1, q_{ij}$	$\log S_0, D_{ij}$	$\text{diag}(s^2)$	0
Diffusion Kurtosis Imaging (DKI) (weighted least-squares) (Veraart et al., 2013)	$\log S(\mathbf{q})$	$1, q_{ij}, q_{ij}q_{kl}q_{ij}$	$\log S_0, D_{ij}, K_{ijkl}$	$\text{diag}(s^2)$	0
Q-space Trajectory Imaging (QTI) (Westin et al., 2016)	$\log S(B)$	$1, B, B^{\otimes 2}$	$\log S_0, \langle D \rangle, \mathbb{C}$	$\text{diag}(s^2)$	0
Constrained Spherical Deconvolution (CSD) (Tournier et al., 2007)	$S(\hat{\mathbf{q}})$	$\chi(\hat{\mathbf{q}})Y_l(\hat{\mathbf{q}})$	c_l	I	L
Q-Ball Imaging (QBI) (Descoteaux et al., 2007; Tuch, 2004)	$S(\hat{\mathbf{q}})$	$Y_l(\hat{\mathbf{q}})$	c_l	I	$\int_{S^2} \ \Delta_b S\ ^2 d\Omega$
MAP-MRI with Laplacian regularization (MAPL) (Fick et al., 2016; Özarslan et al., 2013)	$S(\mathbf{q})$	$\Phi_n(u, q)$	c_l	I	$\int_{\mathbb{R}^3} \ \Delta S\ ^2 d\mathbf{q}$
Spherical Polar Fourier (SPF) (Assemlal et al., 2009)	$S(\mathbf{q})$	$R_k(q)Y_l(\hat{\mathbf{q}})$	c_{kl}	I	$\Lambda_R + \Lambda_Y$

y : response variable

$\mathbf{x}$: input variables

$\phi(\mathbf{x})$: basis function

$\mathbf{c}$: coefficients

$\mathbf{W}$: inverse noise correlation matrix

Λ : regularization matrix

S : signal

S_0 : non-diffusion weighted signal

χ : single fiber response function

Y_l : real spherical harmonics, cf. (Descoteaux et al., 2007)

L : matrix determined iteratively using a sparsifying heuristic

Δ : Laplace operator

Δ_b : Laplace-Beltrami operator

Φ_n : Hermite functions scaled by a factor u

R_k : Gaussian Laguerre polynomials

Λ_R : diagonal matrix penalizing higher radial orders

Λ_Y : diagonal matrix penalizing higher angular orders