Integration of steady-state diffusion MRI with Neural Posterior Estimation (NPE) for post-mortem investigations

Benjamin C. Tendler[1*]
[1]Centre for Integrative Neuroimaging, FMRIB, Nuffield Department of Clinical Neurosciences, University of Oxford, Oxford, UK

*Corresponding Author. Email: benjamin.tendler@ndcn.ox.ac.uk

Abstract

Post-mortem diffusion MRI plays a key role in investigative pipelines to characterise tissue microstructure, with long scan times facilitating the acquisition of datasets with improved spatial/angular resolution and reduced artefacts versus in vivo. Diffusion-weighted steady-state free precession (DW-SSFP) has emerged as a powerful technique for post-mortem imaging, achieving high SNR-efficiency and strong diffusion weighting in the challenging imaging environment of fixed tissue. However, the sophisticated signal forming mechanisms of DW-SSFP limit the integration of advanced microstructural models (e.g. incorporating time-dependence; Monte-Carlo simulations) with parameter estimation routines. Here, I investigate the integration of DW-SSFP with neural posterior estimation (NPE), a parameter inference technique leveraging concepts from Bayesian statistics and machine learning to directly estimate $P(\theta \mid S)$ (i.e. the posterior distribution of parameters θ given signal S). A key challenge is that diffusion attenuation in DW-SSFP is dependent on tissue relaxation properties (T_1/T_2) and transmit inhomogeneity (B_1), which must be incorporated into the NPE network for accurate modelling. By using NPE to estimate $P(\theta \mid S, T_1, T_2, B_1)$ (i.e. conditioning on S and known T_1/T_2/B_1), using a Tensor representation, I demonstrate that NPE achieves accurate parameter estimation even in the presence of non-Gaussian (Rician) noise in low-SNR regimes. Comparisons with conventional non-linear least-squares (NLLS) using both synthetic and experimental DW-SSFP data (whole human post-mortem brain) give excellent agreement, with NPE providing 1000s of posterior samples in a matched evaluation time. Taken together, findings provide a framework to integrate advanced microstructural models with DW-SSFP, and an intuitive approach to incorporate conditional dependencies with NPE.

Highlights

- NPE provides fast and accurate parameter inference for DW-SSFP investigations
- Implemented NPE network accounts for DW-SSFP signal dependencies on T_1, T_2 and B_1
- NPE demonstrates excellent agreement with NLLS using a Tensor representation
- NPE achieves accurate parameter estimation in low SNR non-Gaussian noise regimes
- Evaluations using synthetic & experimental (whole human post-mortem brain) data

Introduction

Non-invasive imaging is a core technology for assessing brain health. Many neuropathologies are characterised by changes at the cellular (μm) scale, motivating the development of imaging techniques sensitive to tissue microstructure. Diffusion MRI is a leading modality in this space, producing images with contrast reflecting the diffusion of water within and surrounding cells (Szafer et al., 1995). Sequence parameters can be adjusted to investigate tissue diffusivity along different spatial orientations (Chenevert et al., 1990; Moseley et al., 1990), length scales, and timing regimes (Bihan, 1995; Vangelderen et al., 1994), with the integration of measured signals with sophisticated biophysical models facilitating the quantitative evaluation of microstructural properties (Alexander et al., 2019; Bihan, 1995; Novikov et al., 2019).

Post-mortem diffusion MRI plays a key role in investigative pipelines to characterise tissue microstructure. Specifically, the long scan times afforded by post-mortem investigations facilitates the acquisition of datasets with improved spatial/angular/b-value resolution and reduced artefacts (e.g. arising from physiological motion) versus in vivo (Roebroeck et al., 2019; Schilling et al., 2025). Post-mortem MRI can also be combined with ultra-high resolution microscopy techniques (performed in the same tissue) to validate microstructural estimates (Choe et al., 2012; Mollink et al., 2017), and directly compared with in vivo MRI to facilitate cross scale-comparisons. More recent work has extended such approaches beyond validation, integrating post-mortem diffusion MRI and microscopy data to address degeneracies (Howard et al., 2019) and improve estimation of brain connectivity (S. Zhu et al., 2025).

A key challenge for post-mortem diffusion MRI investigations is the acquisition of high-quality datasets. Direct translation of existing diffusion MRI protocols can lead to images with poor SNR and low diffusion contrast, arising from the short T_2 and low diffusivity of fixed post-mortem tissue (Roebroeck et al., 2019). Whilst this can be addressed using specialised MRI scanners with powerful gradients, such approaches limit sample size (e.g. pre-clinical) or scanner availability (e.g. CONNECTOM systems) (McNab et al., 2013), motivating investment into alternative diffusion MRI sequences that may be more suitable for post-mortem imaging.

Diffusion-weighted steady-state free precession (Kaiser et al., 1974; Le Bihan, 1988; Merboldt et al., 1989b, 1989a) (DW-SSFP) (Figure 1) has emerged as a powerful technique for post-mortem investigations, with signal-forming mechanisms leading to high SNR-efficiency, strong diffusion-weighting, and minimal geometric-distortions (Miller et al., 2012) (see Theory). SNR-efficiency has been demonstrated to increase with magnetic field strength (Foxley et al., 2014), with recent post-mortem DW-SSFP examples including a 500 μm whole human brain (7 T) (Tendler et al., 2022), a 400 μm whole macaque brain (10.5 T) (Tendler et al., 2025; Warrington et al., 2025), and diffusion MRI across a range of human (Cardenas et al., 2017; Edlow et al., 2018; Ji et al., 2024; Nolan et al., 2021; van Veluw et al., 2019; Wilkinson et al., 2016; Q. Zhu et al., 2024) and non-human (Berns et al., 2015; Berns & Ashwell, 2017; Boch et al., 2024; Bryant et al., 2021; Cook & Berns, 2022; Flem et al., 2025; Orekhova et al., 2022; Roumazeilles et al., 2020; Zheng et al., 2024) brain samples.

DW-SSFP signal measurements have demonstrated distinct microstructural sensitivities versus the diffusion-weighted spin-echo (DW-SE) sequence (McNab & Miller, 2008; Tendler, 2025; Zheng et al., 2024). Whilst previous work has successfully integrated time-independent systems (characterised by Gaussian compartments) with analytical DW-SSFP signal representations for parameter estimation (Foxley et al., 2014; McNab et al., 2009; Tendler, 2025; Tendler, Foxley, Cottaar, et al., 2020), the sophisticated signal-forming mechanisms of DW-SSFP (Freed et al., 2001; Tendler, 2025) currently limits the integration of advanced biophysical models (e.g. incorporating time-dependence or derived from Monte-Carlo simulations) (Novikov et al., 2019) with parameter estimation routines (see Theory). Such extensions would greatly benefit from new data-fitting approaches to support rapid parameter estimation.

In this work, I investigate the integration of DW-SSFP with Neural Posterior Estimation (NPE), a parameter inference technique leveraging concepts from Bayesian statistics and machine learning (Papamakarios & Murray, 2016) to directly estimate the posterior distribution of parameters, $P(\theta \mid S)$ (i.e. the probability distribution over parameters θ given data S). Briefly, given a prior distribution of parameters, $P(\theta)$, and a forward simulation model, $f(\theta \rightarrow S)$, NPE uses simulated data pairs $[\theta, S]$ to train a neural network to estimate $P(\theta \mid S)$. Once trained, experimental data can be passed to the network to estimate $P\left(\theta \mid S_{\mathrm{exp}}\right)$. In comparison to more conventional parameter estimation methods (e.g. non-linear least squares), the estimated posterior

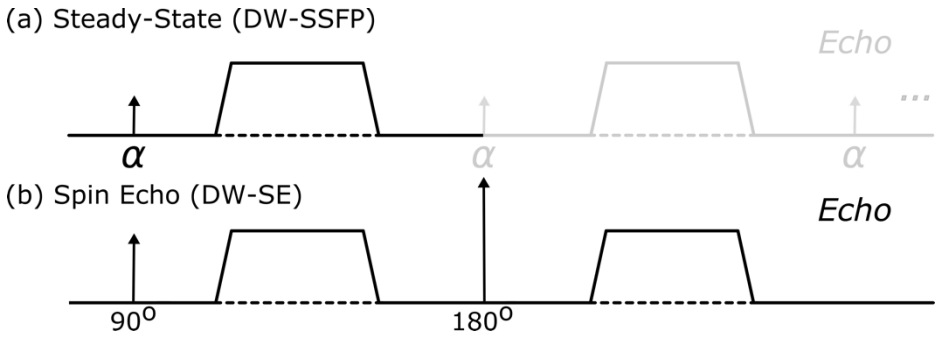

Figure 1: Diffusion encodings. (a) The diffusion encoding of the DW-SSFP sequence consists of a single RF pulse (flip angle α) and diffusion gradient per TR (black line). A short TR ($\sim 20 - 40$ ms) combined with no spoiling of transverse magnetisation leads to magnetisation being sensitised to diffusion gradients across multiple TRs (grey line). This sensitisation leads to magnetisation dephasing and rephasing consistent with the diffusion encoding of conventional diffusion MRI sequences, e.g. the DW-SE displayed in (b), leading to diffusion-weighted contrast.

distribution can provide detailed information about parameter uncertainties and model degeneracies. A key feature of the implemented NPE network is that inference is *amortised*, i.e. training is performed only once to establish a network capable of performing parameter estimation across voxels in a dataset (see Theory). NPE can be considered a specific case of simulation based inference (SBI), which describes a range of approaches to estimate a target distribution using simulated data (Cranmer et al., 2020).

Previous work has demonstrated the potential of NPE for diffusion-weighted spin-echo (DW-SE) investigations (Eggl & De Santis, 2024; Jallais & Palombo, 2024; Manzano-Patron et al., 2025), providing rapid parameter estimates and detailed posterior distributions to characterise uncertainty and model degeneracies. A key challenge for DW-SSFP is that the signal model has additional voxelwise dependencies on tissue relaxation properties (T_1/T_2) and B_1, which must be estimated for accurate diffusion modelling (see Theory). From the perspective of NPE, this corresponds to a different relationship between a sample's diffusivity properties and the measured DW-SSFP signal per T_1/T_2/B_1 combination. Training a different network per T_1/T_2/B_1 combination is infeasible, requiring the adoption of an alternative approach.

Here, I address signal dependences in DW-SSFP by using NPE to estimate $P(\theta \mid S, T_1, T_2, B_1)$, i.e. the posterior distribution conditioned on the measured signal and known T_1, T_2, and B_1 values. Using a Tensor representation (Basser et al., 1994) of the DW-SSFP signal, I demonstrate that the proposed network successfully accounts for DW-SSFP signal dependencies, achieving excellent agreement with conventional non-linear least-squares (NLLS) whilst providing 1000s of posterior samples in a matched evaluation time. By additionally incorporating noisy inputs into the training dataset, I demonstrate that the NPE network can accurately perform parameter estimation in low SNR regimes characterised by non-Gaussian (i.e. Rician) noise distributions.

The presented work aims to improve the characterisation of tissue microstructure using steady-state diffusion MRI sequences. The implemented approach can be adapted to incorporate more sophisticated microstructural models (including Monte Carlo simulations), alternative acquisition schemes, or sequence designs, requiring only a forward simulation model to transform parameters into signal estimates. Experimental evaluations are performed in a DW-SSFP dataset acquired in a whole post-mortem human brain (0.85 mm isotropic resolution). Software, including scripts to replicate many of the findings in this manuscript, is available at github.com/BenjaminTendler/SBI_DWSSFP.

Theory

Overview of Diffusion-Weighted Steady-State Free Precession (DW-SSFP)

The diffusion encoding of the DW-SSFP sequence (Kaiser et al., 1974; Le Bihan, 1988; Merboldt et al., 1989b, 1989a) consists of a single RF pulse and diffusion gradient per TR (Figure 1a). A short repetition time (TR) (~20 – 40 ms) combined with no transverse spoiling leads to magnetisation that experiences repeated sensitisation to diffusion gradient pairs over multiple TRs (Figure 1a – grey line), consistent with the dephasing/rephasing of magnetisation associated with the DW-SE sequence (Figure 1b).

From the perspective of extended phase graphs (EPG) (Weigel, 2015), the steady-state DW-SSFP signal can be interpreted as a composite measurement corresponding to the weighted-sum of magnetisation pathways associated with different degrees of diffusion-weighting (Weigel et al., 2010). Whilst this leads to a comparatively complicated set of signal-forming mechanisms versus the DW-SE, these same signal forming mechanisms offer several advantageous properties for diffusion MRI investigations:

- **SNR-efficiency**: A short TR (~20 – 40 ms) combined with a short diffusion gradient (~10 ms) leads to a large fraction of each TR dedicated to signal readout and a high SNR-efficiency (Miller et al., 2012).

- **Diffusion-weighting**: Many steady-state magnetisation pathways are associated with long diffusion times or experience repeat sensitisation to diffusion gradients across several TRs, leading to a signal with a high 'effective' b-value (Buxton, 1993; Tendler, 2025).

- **Signal sampling**: Sequence timings facilitate a short echo time (TE), advantageous for improving SNR in short T_2 tissue (Foxley et al., 2014; Miller et al., 2004).

- **Geometric distortion**: The sequence's short TR is compatible with a single-line readout, achieving low/negligible geometric distortion.

Whilst routine in vivo investigations are currently limited by the sequences severe motion sensitivity (Merboldt et al., 1989b) (an area of continued investigation) (R. O'Halloran et al., 2015; R. L. O'Halloran et al., 2013; Tendler, 2023), the sequence has become an established method for post-mortem investigations, alongside several 'low-motion' in vivo domains (e.g. cartilage, peripheral nerves) (Bieri et al., 2012; Miller et al., 2004; Z. W. Zhang et al., 2008). For an overview of the properties of DW-SSFP, I recommend the review article by McNab & Miller (McNab & Miller, 2010). For a detailed insight into the signal-forming mechanisms of DW-SSFP, I recommend the article by Tendler (Tendler, 2025).

DW-SSFP Signal Dependencies

A key feature of the DW-SSFP sequence is that diffusion-weighting is dependent on several sequence parameters (diffusion gradient amplitude G; diffusion gradient duration δ, repetition time TR, flip angle α), transmit B_1 (spatially modulating the flip angle α), and sample properties complementary to diffusivity (T_1 and T_2). From the perspective of EPG, these dependencies can be interpreted as modifying the relative weighting of magnetisation pathways with different degrees of diffusion weighting (Weigel et al., 2010).

Figure 2 motivates this concept visually, displaying two DW-SSFP signals with identical profiles, but distinct T_1, T_2, B_1, and diffusivity properties. To accurately estimate diffusion coefficients, these dependencies must be precisely accounted for when performing signal modelling. For dependencies that vary spatially across a dataset (T_1, T_2 and B_1), this requires the acquisition of complementary experimental MRI data using quantitative T_1, T_2 and B_1 mapping techniques.

DW-SSFP Signal Modelling

The complexity of the signal-forming mechanisms of DW-SSFP is reflected in Appendix 1, which presents an analytical expression for the DW-SSFP signal incorporating free Gaussian diffusion (Freed et al., 2001). In contrast to a simple mono-exponential ($S_0 \cdot e^{-bD}$) signal expression for the DW-SE sequence, the DW-SSFP sequence is described using a detailed recursive expression, requiring a dozen equations to represent.

Whilst more computationally demanding, the expression in Appendix 1 is not a barrier itself to the integration of parameter estimation routines with DW-SSFP data. Previous studies have successfully incorporated a (1) Tensor (Foxley et al., 2014), (2) Ball & Sticks (McNab et al., 2009), and (3) Gamma distribution of diffusivities (Tendler, Foxley, Cottaar, et al., 2020) to characterise tissue properties. More recent work introduced an alternative representation of the DW-SSFP signal to facilitate the incorporation of biophysical models (Tendler, 2025), establishing a data fitting approach to obtain NODDI (H. Zhang et al., 2012) estimates from experimental DW-SSFP data. To date, parameter estimation has been perform using iterative NLLS or Markov Chain Monte Carlo (MCMC) methods.

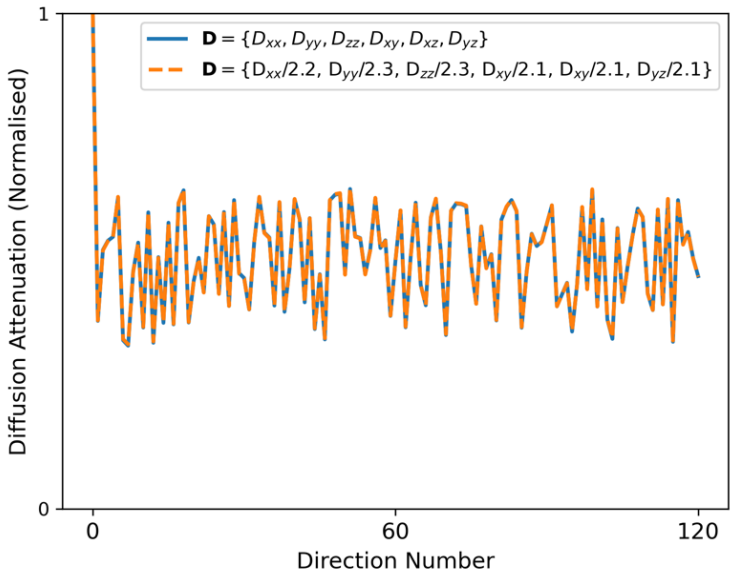

Figure 2: DW-SSFP signal degeneracy. Using a Tensor representation of diffusivities, here I display two simulated DW-SSFP signals with distinct diffusion coefficients but identical diffusion attenuation profiles. Profile degeneracy arises from the dependencies of the DW-SSFP signal on T_1, T_2 and B_1, which additionally modulates the diffusion-weighting of DW-SSFP. Specifically, the two DW-SSFP signals have distinct T_1 (blue: 650 ms; orange: 800 ms), T_2 (blue: 35 ms; orange: 60 ms) and B_1 (blue: 1; orange: 0.5) values, which in combination with differing diffusion coefficients leads to an identical signal profile. Taken together, incorporation of T_1, T_2 and B_1 estimates (using complementary mapping techniques) is required for accurate estimation of diffusion coefficients in DW-SSFP. Simulation performed using the experimental acquisition parameters described in the Methods at a single flip angle ($\alpha = 24°$). Default diffusion coefficients based on Tensor estimates in a single corpus callosum voxel of experimental post-mortem data, defining $\boldsymbol{D} = \{1.63, 0.81, 0.82, -0.41, -0.40, 0.22\} \cdot 10^{-4}$ mm/s.

When considering biophysical models with time-dependent features (e.g. arising from restrictions experienced by mobile spins), forward-simulations can be performed using (1) Monte Carlo simulations or (2) the Gaussian Phase Approximation (Stepišnik, 1993; Tendler, 2025) (GPA). The long times required to perform a single forward simulation using these methods for DW-SSFP limits their integration with iterative parameter estimation routines. Specifically, Monte Carlo simulations require the DW-SSFP signal to reach a steady state over many TRs, increasing simulation time versus an equivalent DW-SE investigation. GPA estimations must be averaged over many signal-forming pathways, requiring extensive parallelisation or limiting parameter estimation to approximation regimes of the DW-SSFP signal (Tendler, 2025). The sequence's dependencies on T_1, T_2 and B_1 additionally limits the availability of alternative dictionary-based fitting methods (requiring three additional dimensions for dictionary construction).

I propose to address this by integrating DW-SSFP with NPE for rapid parameter inference, as described in the next section.

Neural Posterior Estimation (NPE)

NPE leverages advances in machine learning network architectures to obtain detailed information about $P(\theta \mid S)$ (i.e. the probability distribution over parameters θ given data S). In contrast to parameter estimation procedures that provide pointwise estimates of θ, the posterior distributions provided by NPE support the characterisation of parameter estimates, uncertainty, and degeneracies (Cranmer et al., 2020; Jallais & Palombo, 2024). Implemented NPE networks facilitate the rapid and direct characterisation of multidimensional posterior distributions with high sampling density, in contrast to conventional posterior-estimation procedures typically associated with reduced sampling efficiency (e.g. approximate Bayesian computation) or iterative sampling schemes (e.g. MCMC) (Cranmer et al., 2020).

NPE network training is performed using simulated data, synthesising many parameter/data pairs ([θ, S]) estimated from a prior distribution, $P(\theta)$, and a forward simulation model, $f(\theta \rightarrow S)$. Once trained, experimental data (S_{exp}) can be passed to the network to directly estimate $P(\theta \mid S_{\text{exp}})$. Whilst the training

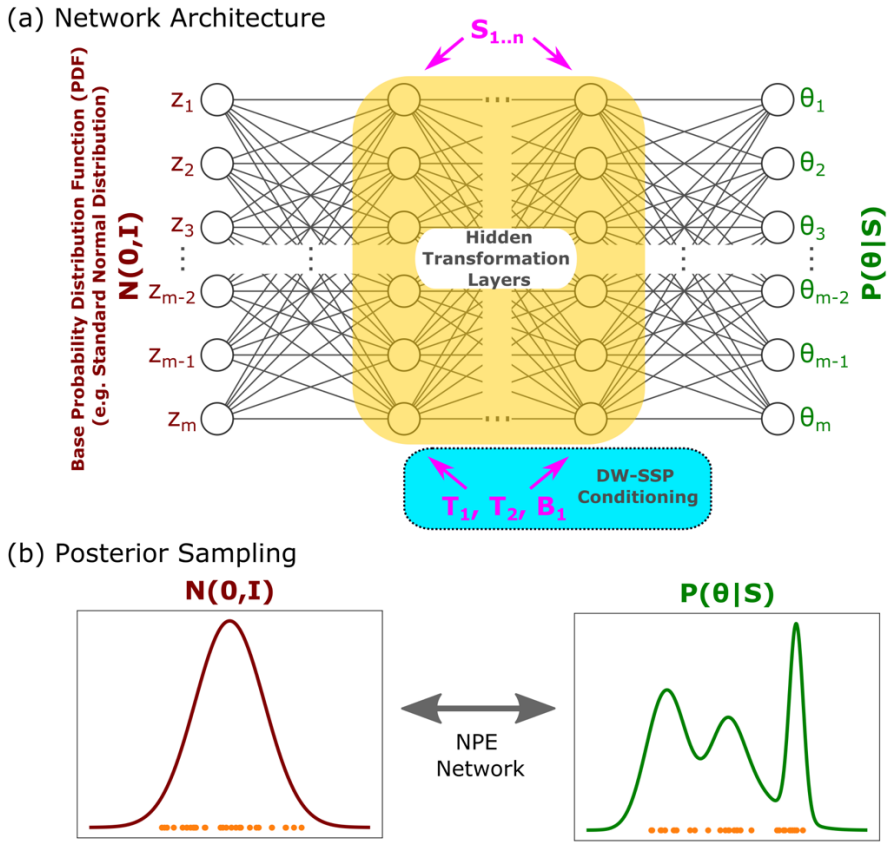

Figure 3: NPE network architecture incorporating DW-SSFP signal dependencies. (a) High-level overview of the network architecture in NPE implemented with normalising flows (Papamakarios et al., 2021), an invertible network learning the transformation between a base PDF (left), and $P(\theta \mid S)$ (right). Conditioning is achieved by passing S (dimensions 1 x n) to the 'hidden' transformation network layers (yellow), where I propose additionally passing T_1, T_2, and B_1 for DW-SSFP conditioning (blue). (b) Once trained, posterior samples, $\theta \sim P(\theta \mid S)$, are generated by transforming base samples, $Z \sim N(0, I)$, through the network. Samples of Z and θ visualised in (b) as orange dots, with the density of sampled points reflecting the PDF.

phase is typically slow (e.g. several hours/days), inference can be amortised (i.e. training only needs to be performed once) and sampling the posterior distribution using the trained network is extremely fast (e.g. milliseconds per voxel).

NPE network architectures are designed to preserve $P(\theta \mid S)$ as a valid probability distribution function (PDF). In this work, I apply NPE using an invertible neural network architecture incorporating normalising flows (Papamakarios et al., 2021) via a neural spline flows (NSF) density estimator (Durkan et al., 2019). Specifically, the network estimates the transformation between a base (e.g. a standard normal distribution, as used in this manuscript (Durkan et al., 2019)) PDF and $P(\theta \mid S)$, with Figure 3a providing a high-level overview of the network design. This approach aims to preserve $P(\theta \mid S)$ as a valid PDF by maintaining key properties of the base PDF as it is transformed through the network (e.g. preserving volume changes), whilst providing a sophisticated transformation approach capable of characterising detailed posterior distributions (Durkan et al., 2019; Papamakarios et al., 2021).

Once trained, samples from the posterior, $\theta \sim P(\theta \mid S)$, can be generated by sampling from the base PDF, $Z \sim N(0, I)$ and passing the samples through the network (Figure 3b). Conditioning is achieved by additionally passing S to the network (i.e. via 'hidden' transformation network layers) (Figure 3a). NPE methods have been shown to accurately approximate sophisticated PDFs (Durkan et al., 2019), with previous work with the DW-SE sequence demonstrating accurate parameter inference for a range of microstructural models (Eggl & De Santis, 2024; Jallais & Palombo, 2024; Manzano-Patron et al., 2025). For further details on normalising flows and NPE, I recommend the articles by Papamakarios et al. (Papamakarios et al., 2021), and Radev et al. (Radev et al., 2023).

DW-SSFP		Turbo inversion-recovery (TIR) – T_1	
q-value (cm^{-1})	300	Resolution (mm^3)	0.9x0.9x0.9
Diffusion Gradient Duration (ms)	13.56	Number of inversions	4
Diffusion Gradient Strength (mTm^{-1})	52	TE (ms)	14.2
Flip angles (°)	24 and 94	TR (ms)	1000
No. directions (per flip angle)	120	TIs (ms)	120, 240, 480 & 935
No. non-DW (per flip angle)	6 (q=20 cm^{-1})	Flip angle (°)	180
Resolution (mm^3)	0.85x0.85x0.85	Turbo Factor	4
TE (ms)	21	Bandwidth (Hz per pixel)	130
TR (ms)	28	Acquisition time (per TI)	40:49
EPI factor	3	Number of averages	1
Bandwidth (Hz per pixel)	393		
Acquisition time (per volume)	5:47	Turbo spin-echo (TSE) - T_2	
Acquisition time (per flip angle)	12:08:42	Resolution (mm^3)	0.9x0.9x0.9
No. of averages	1	Number of echoes	6
		TEs (ms)	12.6:12.6:75.6
Actual flip-angle imaging (AFI) – B_1		TR (ms)	1000
Resolution (mm^3)	3x3x3	Flip angle (°)	180
TE (ms)	1.5	Turbo Factor	6
TR_1/TR_2 (ms)	4.4/11	Bandwidth (Hz per pixel)	166
Flip angle (°)	60	Acquisition time (per TE)	36:01
Bandwidth (Hz per pixel)	630	Number of averages	1
Acquisition time	0:41		
Number of averages	1		

Table 1: Experimental acquisition parameters. Acquisition parameters for the DW-SSFP sequence and complementary data acquired for quantitative T_1, T_2 and B_1 mapping. For full details motivating the choice of sequence parameters, see (Tendler, Foxley, Hernandez-Fernandez, et al., 2020).

To account for dependencies of the DW-SSFP signal on T_1, T_2 and B_1 (see above Theory), I propose additionally passing T_1, T_2 and B_1 to the hidden transformation network layers for conditioning (Figure 3a - blue box). This approach approximates the posterior PDF as $P(\theta \mid S, T_1, T_2, B_1)$, avoiding the requirement to train a separate network per instance of T_1, T_2 and B_1.

Methods:

DW-SSFP forward model

A forward model of the DW-SSFP signal under a Tensor representation was established incorporating DW-SSFP acquisition parameters, T_1, T_2, B_1, and Tensor coefficients (Appendix 2), implemented in Python (3.12.8). DW-SSFP sequence parameters were defined based on the experimental DW-SSFP dataset analysed in this manuscript, provided in Table 1. For full details motivating the choice of sequence parameters, see (Tendler, Foxley, Hernandez-Fernandez, et al., 2020).

Network Implementation

All data analysis was performed in Python (3.12.8), with the conditioned NPE network implemented using the SBI toolbox (0.23.3) (Boelts et al., 2024; Tejero-Cantero et al., 2020) as follows:

- *Priors*: Tensor coefficients ($\mathbf{D}$), uniform distributions with limits:
  - [0, 5] $\cdot\ 10^{-4}$ mm/s for diagonal components (D_{xx}, D_{yy}, D_{zz})
  - [-2.5, 2.5] $\cdot\ 10^{-4}$ mm/s for off-diagonal components (D_{xy}, D_{xz}, D_{yz})

- *Data Pairs*: 6,000,000 data pairs, corresponding to 1,000,000 simulations with 5 Rician noise levels + noise free. Simulated signals corresponded to a 1D array of dimensions 1x252 (Table 1), with each simulation associated with an arbitrary:
  - Tensor (sampled from the prior)
  - T_1 (uniform distribution, limits = [300, 1200] ms)
  - T_2 (uniform distribution, limits = [20, 80] ms)

- B_1 (uniform distribution, limits = [0.2, 1.2])
- SNR (Rician noise simulations) (uniform distribution, limits = [2,50]). SNR was defined with respect to the non-diffusion weighted (i.e. q = 20 cm^{-1}) signal amplitude, setting $S_{noise} = |S + \eta|$, where $\eta \sim N(0, \frac{S_{nonDW}}{\mathrm{SNR}})$.

- *Training*: The NPE network was implemented using an NSF density estimator (Durkan et al., 2019), with default SBI toolbox parameters: NPE-C/APT algorithm (Greenberg et al., 2019); no. hidden features = 50, no. transforms = 5; no. spline bins = 10; parameter and data z-score: independent; one training round. T_1, T_2 and B_1 conditioning was achieved by appending their values to the signal during training, creating a 1D array corresponding to [S, T_1, T_2, B_1] of dimensions 1x255.

Similar to the approach in (Manzano-Patron et al., 2025), an additional classifier network was trained to identify regions of the prior distribution (i.e. parameter combinations) that produced positive semi-definite tensors, implemented using the *RestrictionEstimator* function in the SBI toolbox (ResNet; no. hidden features = 100, no. blocks = 2; dropout prob. = 0.5; parameter z-score: independent) (Deistler et al., 2022; Lueckmann et al., 2017). This can be interpreted as modifying the prior distribution prior to synthesising data pairs to avoid invalid parameter combinations. The classifier network was trained using 500,000 simulated data pairs, rejecting 53.7% of parameter combinations, accelerating data synthesis by 134.7%.

The total training time was ~34 hours on a personal laptop (Macbook Pro, Sonoma 14.5, M1, 16GB RAM). This consisted of ~33 hours for the NPE network training and < 1 hour to train the classifier network and synthesise the noisy data pairs. The trained NPE network required 1.1 MB in storage space.

Network Evaluation: Synthetic data

To evaluate the performance of the conditioned NPE network, simulated DW-SSFP signals were passed to the trained network for parameter inference. $\mathbf{D}$ was subsequently estimated from the simulated signals and compared to the known ground truth. Two investigations were performed to estimate parameters for (1) fixed $\mathbf{D}$ across a range of T_1, T_2 and B_1 values, and (2) arbitrary $\mathbf{D}$, T_1, T_2 and B_1 combinations. 100 posterior samples were taken per simulation, with the mean of the posterior estimates used for evaluation.

To characterise the networks robustness to data noise, a similar evaluation was performed as a function of SNR. $\mathbf{D}$ was estimated and compared to the known ground truth from simulated DW-SSFP signals with added Rician noise, alongside an equivalent NLLS fitting implementation with SciPy *curve_fit* (1.12.0) (trust region reflective algorithm) (E. Jones et al., 2001). Two investigations were performed to estimate parameters for (1) fixed $\mathbf{D}$ across a range of SNR levels (5 to 50), and (2) arbitrary $\mathbf{D}$, T_1, T_2 and B_1 combinations in a low SNR regime (SNR = 5). 100 posterior samples were averaged per NPE simulation, with parameter estimates additionally averaged over 1000 noise realisations per SNR level (NPE and NLLS) for the fixed $\mathbf{D}$ investigation.

To assess the speed of NPE for parameter estimation, comparisons were made to NLLS as a function of the number of posterior samples, evaluating 1000 arbitrary $\mathbf{D}$, T_1, T_2 and B_1 combinations with SNR = 20. Finally, the number of samples required to robustly estimate the mean and standard deviation of the Tensor coefficients for the implemented DW-SSFP acquisition scheme (Table 1) was assessed as a function of SNR for 1000 arbitrary $\mathbf{D}$, T_1, T_2 and B_1 combinations.

Network Evaluation: Experimental data

Experimental DW-SSFP and T_1/T_2/B_1 mapping data were acquired in a single whole human post-mortem brain on a 7T Siemens system (32ch-receive/1ch-transmit head coil) using the protocol in Table 1 (Tendler, Foxley, Hernandez-Fernandez, et al., 2020). Prior to fitting, experimental data were pre-processed followed the procedure in (Tendler, Foxley, Hernandez-Fernandez, et al., 2020). Briefly, a quantitative T_1 map was estimated from the TIR data assuming mono-exponential signal recovery. A quantitative T_2 map was estimated from the TSE data using an EPG signal model as previously described (Tendler et al., 2021). A quantitative B_1 map was estimated from the AFI data following the processing steps in the original manuscript (Yarnykh, 2007). The T_1, T_2 and B_1 maps were coregistered to the DW-SSFP data using a 6 degrees of freedom (DOF) transformation with FSL FLIRT (Jenkinson et al., 2002; Jenkinson & Smith, 2001).

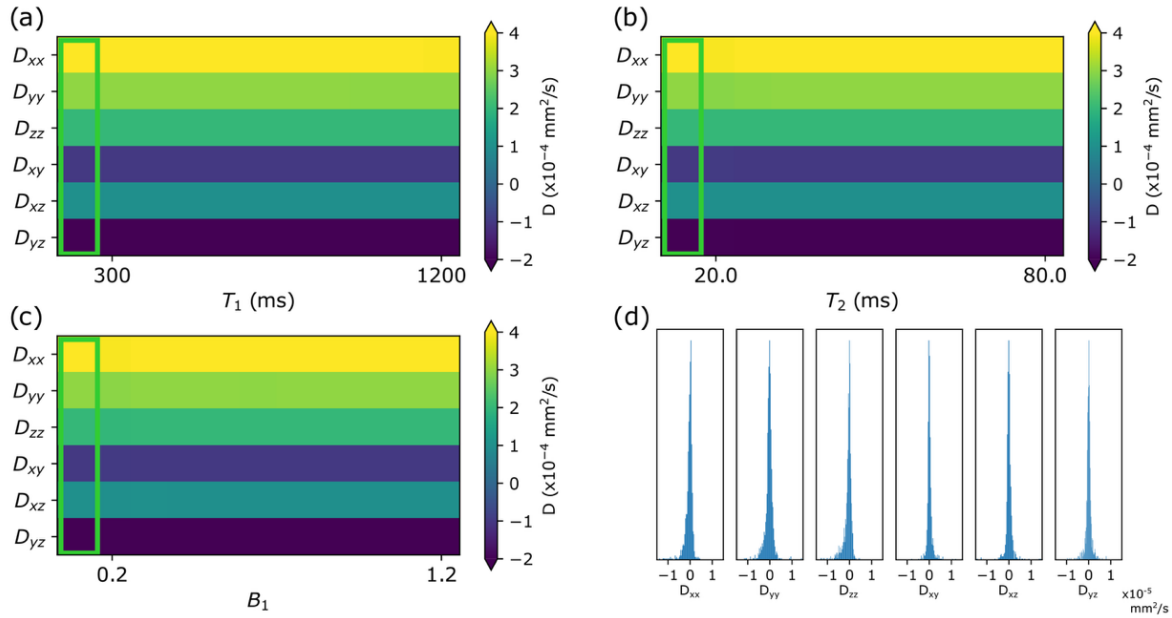

Figure 4: Robustness of implemented NPE network to DW-SSFP signal dependencies. (a-c) display the ground truth (green box) versus estimated Tensor coefficients as a function of (a) T_1, (b) T_2, and (c) B_1, setting $\mathbf{D} = \{4, 3, 2, -1, 1, -2\} \cdot 10^{-4}$ mm/s Excellent agreement is found across the dependency range, with median absolute difference = 0.27% and interquartile range = 0.37%. (d) displays the average difference (plotted as a histogram) between estimated and ground truth Tensor coefficients across 1000 simulations with arbitrary $\mathbf{D}$, T_1, T_2, and B_1 (estimated from the prior - see Methods). Excellent agreement is again found with respect to ground truth (median absolute difference = 5.7 $\cdot 10^{-7}$ mm^2/s; interquartile range = 8.2 $\cdot 10^{-7}$ mm^2/s). 100 posterior samples were taken per simulation, with the mean of the posterior estimates used for comparison. Default values (a-c): T_1 = 600 ms; T_2 = 33.3 ms; B_1 = 0.98, corresponding to grid elements closest to the mean relaxation values estimated in the acquired post-mortem experimental data and no transmit inhomogeneity (T_1 = 650 ms; T_2 = 35 ms; B_1 = 1).

For the DW-SSFP data, a Gibbs ringing correction was first applied (Kellner et al., 2016). Individual volumes were subsequently coregistered using a 6 DOF transformation with FSL FLIRT. The DW-SSFP magnitude data was reconstructed using a root sum-of-squares coil combination, leading to a noncentral chi noise distribution (Aja-Fernández et al., 2009). To reduce parameter bias in low SNR regions (D. K. Jones & Basser, 2004) and facilitate comparisons between NPE and NLLS, a noise floor correction was performed to the experimental data (Gudbjartsson & Patz, 1995b), with the noise floor estimated from the background signal. To circumvent the requirement of estimating the signal amplitude (S_0) as part of the parameter estimation procedure, S_0 was estimated from the non-diffusion weighted (q = 20 cm^{-1}) volumes using the expression in Appendix 1 and removed from the experimental data (via division) prior to fitting.

The DW-SSFP signal, alongside T_1, T_2 and B_1 maps were passed to the NPE network on a voxelwise basis to estimate $P(\mathbf{D} \mid \mathrm{S}, T_1, T_2, B_1)$, with 100 posterior samples estimated per voxel. The mean of the posterior samples were used for comparison to the NLLS fitting implementation. In a small fraction of voxels ($\sim$1.3% of the dataset), Tensor coefficients fell outside of the prior bounds ($> 5 \cdot 10^{-4}$ mm/s) and could not be estimated using NPE. This predominantly corresponded to voxels at the brain surface associated with residual fixative trapped within brain sulci or ventricles, and were not incorporated into the remaining analysis.

Results:

Simulations

Figure 4a-c compares estimated Tensor coefficients as a function (a) T_1, (b) T_2, and (c) B_1 using NPE for fixed $\mathbf{D}$. Excellent agreement is found versus ground truth (median absolute difference = 0.27%; interquartile range = 0.37%) across the range of T_1, T_2, and B_1, demonstrating that the implemented NPE network can accurately account for DW-SSFP signal dependencies. Figure 4d displays histograms of the difference between the estimated and ground truth Tensor coefficients across 1000 simulations incorporating arbitrary $\mathbf{D}$, T_1, T_2, and B_1

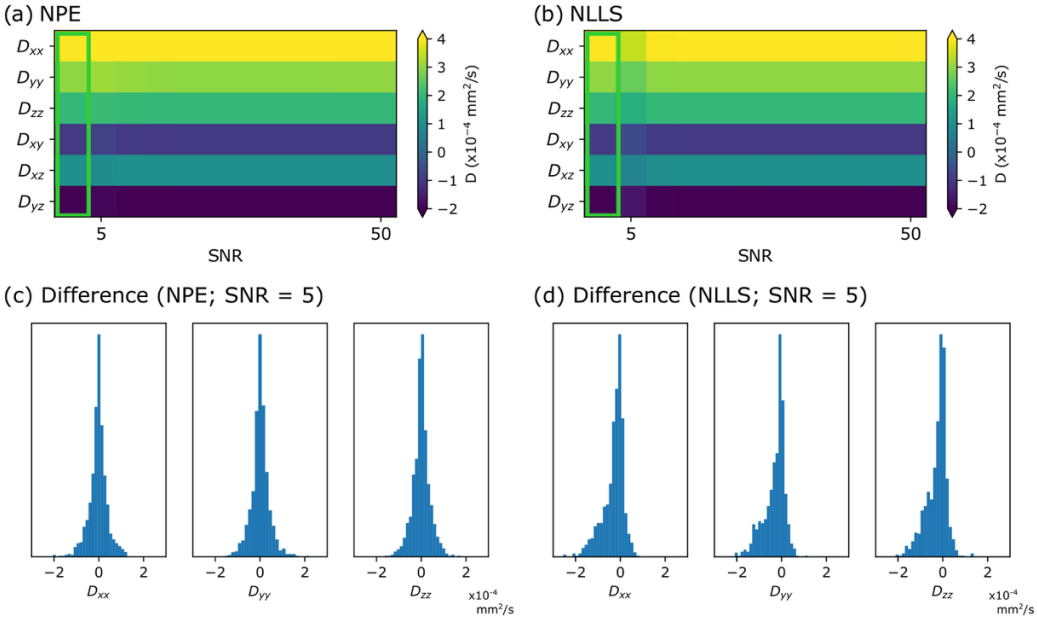

Figure 5: Robustness of implemented NPE network to SNR. Here I display the ground truth (green box) versus estimated Tensor coefficients for the (a) NPE network and (b) NLLS as a function of SNR (Rician noise), setting $\mathbf{D} = \{4, 3, 2, -1, 1, -2\} \cdot 10^{-4}$ mm/s and defining SNR with respect to the non-diffusion weighted signal. Excellent agreement is found across the SNR range (median absolute difference = 0.6%; interquartile range = 1.0%), reflecting the robustness of NPE to parameter estimation within the trained SNR regimes. The NLLS implementation displays increased biases at low SNR values, reflecting the difference between the assumed NLLS noise model (Gaussian) and the underlying noise model of the simulated data (Rician), which are not visible in the NPE estimates. These biases are visualised in (c) and (d), which display the average difference (plotted as a histogram) between estimated and ground truth diagonal Tensor coefficients across 1000 simulations with SNR = 5 and arbitrary $\mathbf{D}$, T_1, T_2, and B_1 (estimated from the prior and uniform distributions - see Methods). Whilst NPE (c) displays symmetric residuals, NLLS (d) is biased towards negative values, indicative of an underestimation of $\mathbf{D}$. Results in (a) and (b) reflect averages from parameter estimates derived from 1000 noise realisations per SNR level. Default values: T_1 = 650 ms; T_2 = 35 ms; and B_1 = 1; corresponding to the mean relaxation values estimated in the acquired post-mortem experimental data and no transmit inhomogeneity.

combinations, again finding excellent agreement (median absolute difference = 5.7 $\cdot 10^{-7}$ mm^2/s; interquartile range = 8.2 $\cdot 10^{-7}$ mm^2/s). Taken together, these findings demonstrate the robustness of NPE to parameter estimation within the prior range.

Figure 5a and b compares the accuracy of parameter estimation for NPE and NLLS across different SNR levels with added Rician noise and fixed $\mathbf{D}$. The NPE network demonstrates excellent agreement to ground truth across the SNR range (median absolute difference = 0.6%; interquartile range = 1.0%), with improved parameter estimation in low SNR regions compared to NLLS. Figure 5c and d displays the residuals (difference from the ground truth) for the diagonal Tensor coefficients estimated with NPE and NLS. Here NPE residuals display a symmetric deviation from ground truth, with the NLLS equivalent biased towards negative values, indicative of an underestimation of Tensor coefficients. These differences reflect the potential of NPE to incorporate the underlying noise distribution (Rician) when performing parameter inference, with NLLS (assuming a Gaussian noise distribution) leading to biased parameter estimates arising from the impact of the noise floor (D. K. Jones & Basser, 2004).

Figure 6 compares the relative time performance of NLLS versus NPE as a function of the number of posterior samples. NPE can sample several thousand posterior estimates in the time required to obtain a single NLLS estimate (54.27 ms). Modelling NPE evaluation time as a linear function (Figure 6b – blue line), I estimate:

$$\mathrm{NPE_{Time}} \approx 2.95\ \mathrm{ms} + 11.45\ \mu\mathrm{s} \cdot \mathrm{N_{Samples}}, \qquad [1]$$

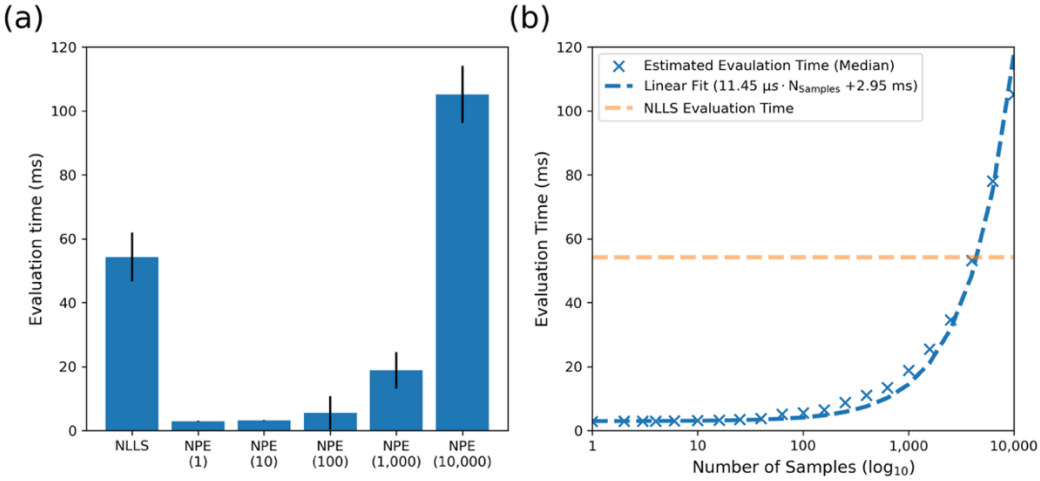

Figure 6: NLLS and NPE evaluation time. Simulating 1000 DW-SSFP signal profiles with arbitrary $\mathbf{D}$, T_1, T_2, and B_1 (estimated from prior and uniform distributions - see Methods) and SNR = 20 (defined relative to the non-diffusion weighted signal), (a) displays the relative evaluation time (median $\pm$ interquartile range) for NLLS and NPE as a function of the number of posterior samples (brackets along x-axis). Here we observe that we can estimate several thousand posterior samples in an equivalent NLLS evaluation time. Investigating this further (b), we can approximate evaluation time in NPE as a linear function (dashed blue line) of the number of samples and an initialisation time (offset), fitting to the median values across the sample range. Using a linear function (Eq. [1]), I estimate ≈ 4500 posterior samples in matched estimation time to NLLS. Note that the x-axis in (b) is $\log_{10}$ scale.

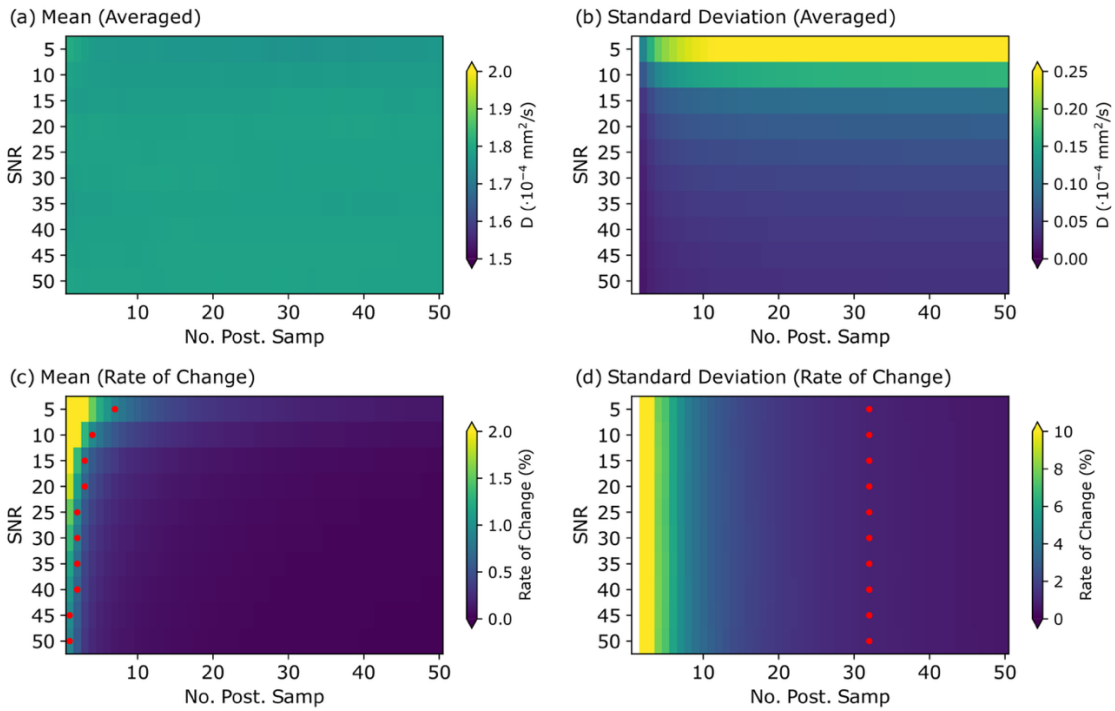

Figure 7: Number of samples required to characterise the posterior. Here I visualise the average (a and b) and rate of change (c and d) of the posterior mean (left column) and standard deviation (right column) as a function of the number of posterior samples (x-axis) and SNR (y-axis). The red dots in (c) and (d) correspond to the number of samples where the rate of change falls to below 1%. Figure calculated from 1000 simulations with arbitrary $\mathbf{D}$, T_1, T_2, and B_1 (estimated from the prior and uniform distributions - see Methods) using the DW-SSFP acquisition scheme described in Table 1, with estimates averaged over the Tensor coefficients and simulations. To improve visualisation and quantification, a smoothing filter (*Savitzky-Golay*) was applied to the averaged data to reduce local fluctuations. Note that the colourmaps in (a)/(b) and (c)/(d) correspond to different ranges.

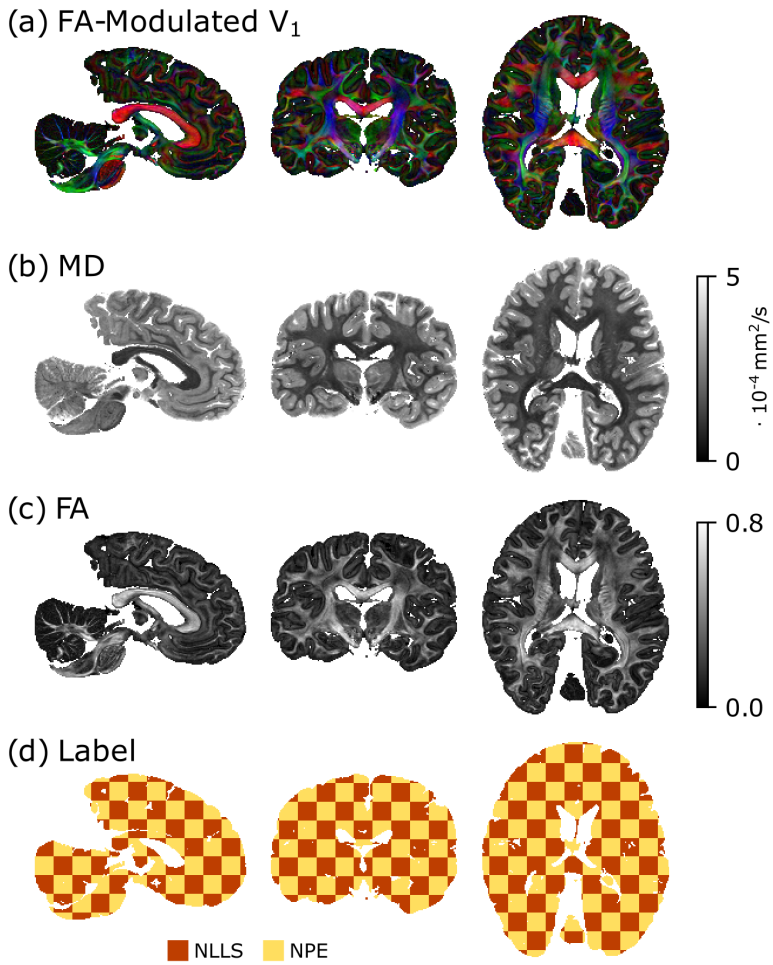

Figure 8: Visual comparison of experimental NPE & NLLS Tensor outputs. Composite chequerboard image (d) combining Tensor outputs from the NPE and NLLS fitting to display (a) $\vec{V}_1$ (modulated by FA), (b) MD and (c) FA maps. Excellent visual agreement is found across the brain, with no visual distinction between the NPE and NLLS outputs. Displayed NPE outputs derived from to the mean of the 100 posterior estimates per voxel.

where $\mathrm{N}_{\mathrm{Samples}}$ is the number of posterior samples. Here we can interpret the offset (2.95 ms) as the initialisation time for NPE, and 11.45 µs as the approximate time required per posterior sample. Using this formula, I estimate ≈ 4500 posterior samples in matched estimation time to NLLS.

For the final simulation investigation, Figure 7 displays the mean and standard deviation of the averaged Tensor coefficients as a function of the number of posterior samples and SNR (based on the DW-SSFP acquisition scheme in Table 1). Specifically, Figures 7a and b display the averaged mean and standard deviation (over tensor coefficients and simulations), with Figures 7c and d displaying the rate of change of the mean and standard deviation as a function of the number of posterior samples. The estimated mean (Figure 7a) is similar across the SNR range (consistent with Figure 5), with the standard deviation (Figure 7b) decreasing with increased SNR, consistent with our expectation of parameter uncertainty. The rate of change analysis (Figures 7c and d) indicates that a few samples (<10) are sufficient to robustly estimate the mean (defined here as <1% change), with ~30 samples required to estimate the standard deviation. Assuming 30 samples corresponds to a 'fully sampled' posterior distribution, NPE provides a 16-17x acceleration versus NLLS based on the DW-SSFP acquisition scheme (Table 1) and Eq. [1].

Experimental Analysis

Figure 8 visualises the experimental Tensor estimates using NPE and NLLS. Excellent visual agreement is found for the derived Tensor outputs, with the composite chequerboard images (label in Figure 8d) displaying no visual distinctions between NPE and NLLS in the reconstructed $\vec{V}_1$ (Figure 8a), MD (Figure 8b) and fractional anisotropy (FA) (Figure 8c) maps.

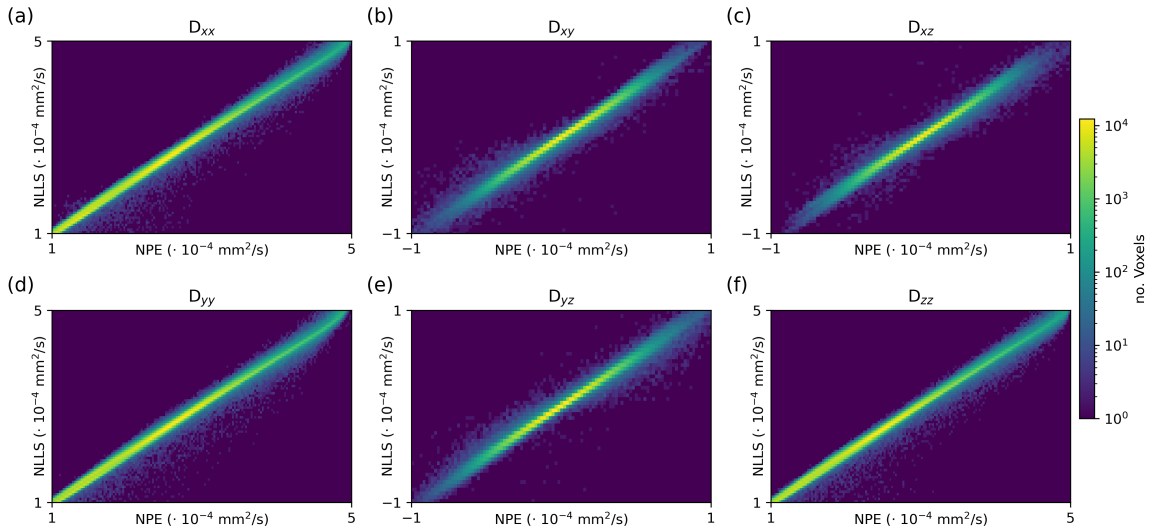

Figure 9: Correlation between NPE & NLLS Tensor elements. 2D histogram distributions correlating individual of NPE (x-axis) and NLLS (y-axis) Tensor elements. Excellent agreement is found across the coefficient distributions, with the histogram density defined predominantly along the diagonal. Pearson correlation coefficients for all elements >0.99. Individual histograms formed on a 200x200 grid, with the colormap reflecting a logarithmic scale (see colour bar).

Figure 9 quantifies the relationship between the NPE and NLLS outputs, displaying their correlation for individual Tensor elements. An strong linear relationship is found, estimating Pearson correlation coefficients r = 0.996 (D_{xx}), 0.993 (D_{xx}), 0.993 (D_{xz}), 0.996 (D_{yy}), 0.994 (D_{yz}), 0.996 (D_{zz}). A small bias (diagonal elements) and asymmetry (off-diagonal elements) in the correlation distribution around the unity line reflects a small underestimation of diffusion coefficients in NLLS versus NPE, likely corresponding to a bias arising from non-Gaussian noise in low SNR regions.

A key additional advantage of NPE versus NLLS is the ability to characterise detailed posterior distributions. Figure 10 displays an example posterior distribution for a single voxel in the corpus callosum, providing detailed information characterising the precision of parameter estimates (diagonal 1D histograms) and relationship between parameters (off-diagonal 2D histograms). Tensor estimates are informative of a corpus callosum voxel, indicating a larger diffusion coefficient along a single axis most closely aligned with the fibre orientation (D_{xx}), with smaller coefficients for orthogonal elements (D_{yy} and D_{zz}). The distribution was formed from 4,500 posterior samples, equivalent to the number of samples that can be achieved with a matched estimation time versus NLLS (Figure 6 and Eq. [1]).

Finally, Figure 11 visualises the average voxelwise signal across diffusion-weighted datasets (Figure 11a) and the corresponding standard deviation of estimated Tensor coefficients (Figure 11b). Regions of reduced signal are associated with higher standard deviations, consistent with the expectations of uncertainty quantification and findings presented in Figure 7. Higher standard deviations are predominantly associated with regions close to the brain boundary along the super-inferior axis (associated with low B_1 in this dataset – see (Tendler, Foxley, Hernandez-Fernandez, et al., 2020)) and deep-grey matter regions (associated with very low T_2 in fixed post-mortem tissue), with the majority of voxels associated with a small standard deviation (order of 10^{-5} mm²/s).

Discussion

Overview

The findings presented here demonstrate that NPE achieves fast and accurate parameter inference for DW-SSFP investigations, with the proposed conditioning approach accounting for dependencies on T_1, T_2 and B_1. Resulting parameter estimates give excellent agreement to ground-truth simulations (Figures 4 and 5) and experimental parameters derived with NLLS (Figures 8 and 9). NPE provides detailed parameter information, characterising full posterior distributions (Figure 10) with extremely fast inference across multidimensional parameter spaces (Figures 6 and 7), quantification of uncertainty (Figure 11), and low computational requirements.

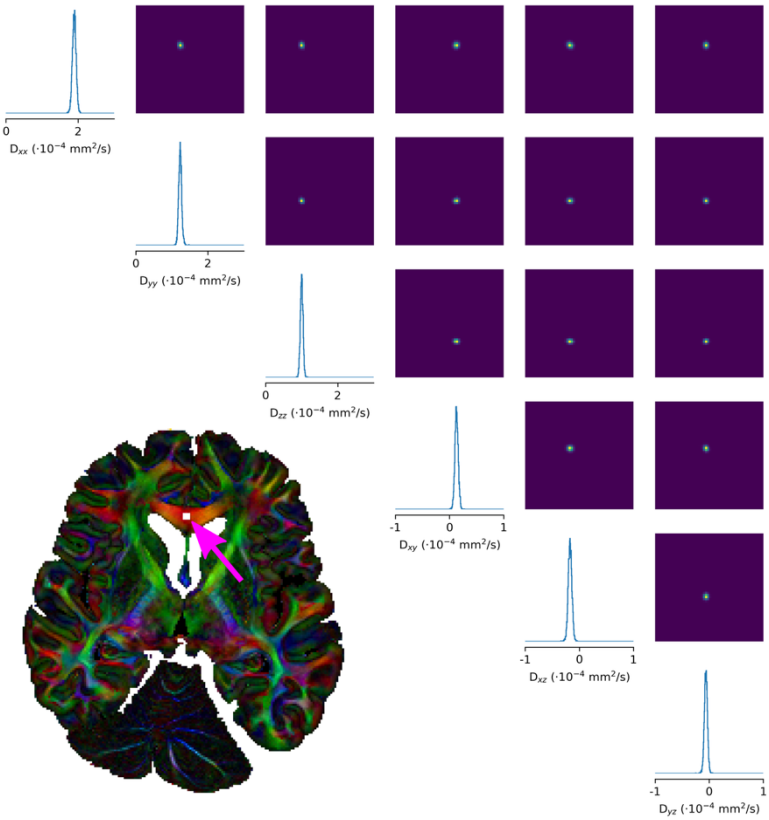

Figure 10: NPE posterior distribution in a single voxel. Here I visualise the NPE posterior distribution for a single experimental voxel on the corpus callosum midline (pink arrow). The diagonal (1D histograms) and off-diagonal (2D histograms) correspond to the univariate and pairwise marginal distributions. The experimental data demonstrates high confidence in the parameter estimates, with a narrow distribution of posterior samples and no degeneracy (i.e. a single peak). Posterior distribution formed from 4500 samples, equivalent to a matched estimation time versus NLLS (Eq. [1]).

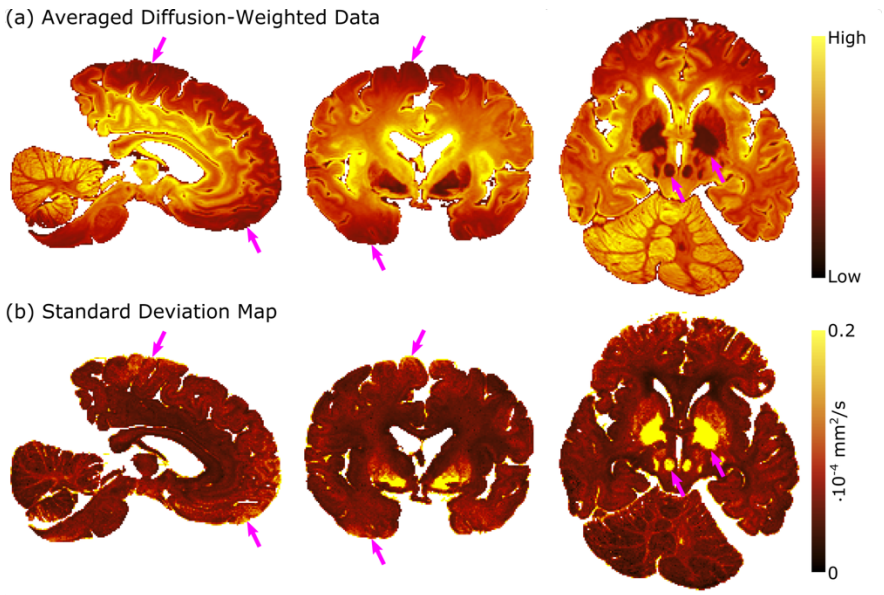

Figure 11: SNR and standard deviation of Tensor coefficient estimates. (a) displays the average signal amplitude across all diffusion-weighted volumes (Table 1), a proxy for SNR. (b) displays the voxelwise standard deviation of Tensor coefficients averaged over all six estimated tensor components. Voxels with a reduced signal amplitude are typically associated with a higher standard deviation, with the pink arrows highlighting example regions. The standard deviation map in (b) provides a quantitative estimate of uncertainty in parameter estimates for NPE, here defined as the average parameter standard deviation from the 100 posterior samples per voxel. Uncertainty is low (order of 10^{-5} mm^2/s) relative to the average amplitude of diffusion coefficients per voxel (e.g. MD map in Figure 8b), with the exception of deep grey matter regions characterised by very low SNR in the acquired post-mortem data.

SNR and noise distributions

It is well established that MRI magnitude data reconstructed from single or multiple channel coil arrays have non-Gaussian noise (typically characterised by e.g. Rician or noncentral chi distributions) (Aja-Fernández et al., 2009; Gudbjartsson & Patz, 1995a). Figure 5 demonstrates that NPE can successfully account for non-Gaussian noise distributions across a range of SNR levels, achieving unbiased parameter estimates when trained on simulated data with added Rician noise.

When considering different noise distributions (e.g. theoretically calculated or estimated from experimental data), these could be similarly incorporated into the forward simulation model for NPE training. Alternatively, NPE networks could be conditioned on noise floor estimates via $P(\theta \mid \mathrm{S}, T_1, T_2, B_1, \mathrm{S}_{\mathrm{nf}})$, incorporated into the forward simulation model via $S_{\mathrm{measured}} = (S^2 + \mathrm{S}_{\mathrm{nf}}^2)^{1/2}$ (Gudbjartsson & Patz, 1995b). Experimental noise floor estimates (i.e. calculated from the background signal) could then be passed to the network during parameter inference to reduce parameter biases.

Deviations from Tensor representation

Whilst the synthetic training data precisely characterises the DW-SSFP signal under a Tensor representation, it is well established that experimental diffusion MRI data can be characterised by a more sophisticated orientation function. One example of this is in crossing fibre regions, characterised by a low FA when fitting a single Tensor to experimental data. This raises a question of whether NPE will fail when performing parameter estimation in tissue regions that have signal representations that deviate substantially from the training data, or whether the detailed posterior distributions will demonstrate sensitivity to multiple fibre populations.

To investigate whether such effects were present in the current dataset, NPE Tensor estimates were investigated in the centrum semiovale (Figure 12), a white matter crossing fibre region. In both NPE and NLLS this region was characterised by a low FA (Figure 12a and b), with no evidence of NPE distinguishing individual fibre populations in the posterior distribution (Figure 12c). Specifically, NPE consistently identified a single peak, with estimated diffusion coefficients characteristic of isotropic diffusion (consistent with NLLS). Whilst previous work has demonstrated multimodal posterior distributions under a Ball and Sticks model (Manzano-Patron et al., 2025), one interpretation is that the underlying Tensor fit to experimental data is not truly degenerate (i.e. characterising diffusivity as more isotropic is a better representation of the experimental data versus accurately characterising one of the Tensors), which can be effectively captured by NPE.

Limitations

One limitation of NPE for diffusion MRI investigations is that networks typically incorporate a specific forward model, defined by a fixed number of parameters or signal estimates. Whilst changes in the diffusion MRI acquisition scheme (e.g. gradient orientations, number of directions, number of shells) can be trivially incorporated with conventional iterative fitting methods (e.g. NLLS, MCMC), in NPE these changes can require a different network architecture. At one extreme, this requires full retraining of the network to account for sequence parameter changes. Such challenges can arise even in the circumstance of missing data (e.g. arising from corrupted diffusion MR volumes).

When considering missing data, imputation methods can be introduced to approximate the missing information. This can be performed by the user (e.g. interpolation to approximate the signal arising from a missing diffusion MRI volume), or leveraging advances in SBI to jointly optimise a data imputation and NPE network (Verma et al., 2025). When considering changes in b-vectors (i.e. orientation or number), dimensionality reduction techniques can also be utilised to condense the acquired data into a subset of orientation-invariant data 'features' prior to NPE. These can be similarly defined by the user (e.g. spherical harmonics) (Descoteaux et al., 2007; Manzano-Patron et al., 2025) or using a joint dimensionality reduction and NPE network architectures (Radev et al., 2023). Addressing such challenges is an area of active investigation in NPE.

The proposed conditioning method introduced in this manuscript may similarly facilitate the generalisability of NPE networks for diffusion MRI investigations. Specifically, NPE networks could be trained over a range of signals arising from different sequence parameters combinations, with the sequence parameters provided to

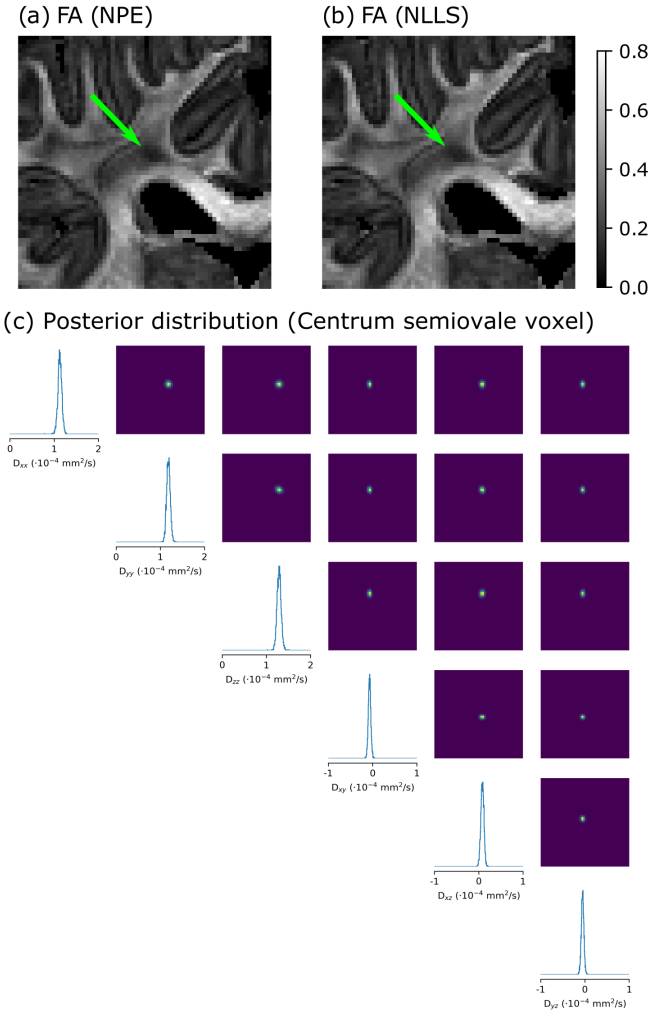

Figure 12: Comparisons of FA estimates in the centrum semiovale. Here I visualise NPE (a) and NLLS (b) estimates of FA in the centrum semiovale of the post-mortem brain, a white matter crossing fibre region. Both NPE and NLLS characterised the region with low FA (green arrow). Posterior distributions estimated from NPE, displayed in (c) for a single centrum semiovale voxel with low FA, were characteristic of isotropic diffusion, with similar diffusion coefficients across the diagonal elements (D_{xx}, D_{yy}, D_{zz}) and low diffusivity in the off-diagonal elements (D_{xy}, D_{xz}, D_{yz}). (c) formed from 4,500 posterior samples.

the NPE network for conditioning. For example, considering a DW-SE investigation, an NPE network could be conditioned on the b-value via $P(\theta \mid S, b)$, circumventing the requirement of explicitly training a different network per b-value. Such conditioning could be combined with missing data and dimensionality reduction techniques to improve the generalisability of NPE networks across acquisition schemes and datasets.

Future directions

Whilst the work presented here aims to demonstrate the feasibility of NPE for DW-SSFP investigations by utilising a signal representation that can be evaluated using complementary methods (e.g. Tensor and NLLS), the implemented routine can be extended to incorporate more sophisticated microstructural models, including Monte Carlo simulations and alternative DW-SSFP signal representations that incorporate time-dependence (Tendler, 2025). Whilst a Tensor representation has a relatively simple posterior distribution (typically characterised by a single symmetric peak in a multidimensional space), previous work has demonstrated that NPE can incorporate sophisticated posterior distributions, with demonstrated applications for the DW-SE sequence (Jallais & Palombo, 2024; Manzano-Patron et al., 2025). The work presented here provides a natural extension for DW-SSP investigations, or alternative steady-state diffusion MRI sequences, including fingerprinting methods.

Conclusion

NPE demonstrates fast and accurate parameter inference for DW-SSFP investigations. The proposed method provides an intuitive approach to incorporate conditional dependencies with NPE, successfully accounting for dependencies of the DW-SSFP signal on relaxation times and transmit inhomogeneity. Presented findings provide a framework for integrating advanced microstructural models with steady-state diffusion MRI sequences, generating detailed posterior distributions for the characterisation of parameter estimates, uncertainty and degeneracies.

Data Statement

Software for the NPE implementation, including scripts to replicate many of the findings presented in this manuscript, are available at github.com/BenjaminTendler/SBI_DWSSFP. The original processed post-mortem data associated with this project can be accessed via the Digital Brain Bank (open.win.ox.ac.uk/DigitalBrainBank) (Tendler et al., 2022) (*Human ALS MRI-Histology dataset*).

Ethics

Post-mortem tissue was provided by the Oxford Brain Bank, a research ethics committee (REC) approved, HTA regulated research tissue bank. Research was conducted under the Oxford Brain Bank's generic Research Ethics Committee approval (15/SC/0639).

Acknowledgements

BCT would like to thank Jose Pedro Manzano Patron for helpful discussions during the development of this project and manuscript feedback.

BCT is funded by a Sir Henry Wellcome Postdoctoral Fellowship (Wellcome Trust) [222829/Z/21/Z]. Human tissue was provided by the Oxford Brain Bank, funded by the Medical Research Council, Brains for Dementia Research, and the NIHR Oxford Biomedical Research Centre. Acquisition of the post-mortem dataset was funded by the Medical Research Council [MR/K02213X/1]. The Centre for Integrative Neuroimaging was supported by core funding from the Wellcome Trust [203139/Z/16/Z and 203139/A/16/Z]. This research was funded in whole, or in part, by the Wellcome Trust [222829/Z/21/Z]. For the purpose of open access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

References

Aja-Fernández, S., Tristán-Vega, A., & Alberola-López, C. (2009). Noise estimation in single-and multiple-coil magnetic resonance data based on statistical models. *Magnetic Resonance Imaging*, *27*(10), 1397–1409.

Alexander, D. C., Dyrby, T. B., Nilsson, M., & Zhang, H. (2019). Imaging brain microstructure with diffusion MRI: practicality and applications. *NMR in Biomedicine*, *32*(4), e3841.

Basser, P. J., Mattiello, J., & LeBihan, D. (1994). MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*, *66*(1), 259–267.

Berns, G. S., & Ashwell, K. W. S. (2017). Reconstruction of the cortical maps of the Tasmanian tiger and comparison to the Tasmanian devil. *PLoS ONE*. https://doi.org/10.1371/journal.pone.0168993

Berns, G. S., Cook, P. F., Foxley, S., Jbabdi, S., Miller, K. L., & Marino, L. (2015). Diffusion tensor imaging of dolphin brains reveals direct auditory pathway to temporal lobe. *Proceedings of the Royal Society B: Biological Sciences*. https://doi.org/10.1098/rspb.2015.1203

Bieri, O., Ganter, C., Welsch, G. H., Trattnig, S., Mamisch, T. C., & Scheffler, K. (2012). Fast diffusion-weighted steady state free precession imaging of in vivo knee cartilage. *Magnetic Resonance in Medicine*. https://doi.org/10.1002/mrm.23061

Bihan, D. Le. (1995). Molecular diffusion, tissue microdynamics and microstructure. *NMR in Biomedicine*, *8*(7), 375–386.

Boch, M., Karadachka, K., Loh, K. K., Benn, R. A., Roumazeilles, L., Bertelsen, M. F., Manger, P. R., Wriggelsworth, E., Spiro, S., Spocter, M. A., Johnson, P. J., Avelino-de-Souza, K., Patzke, N., Lamm, C., Miller, K. L., Sallet, J., Khrapitchev, A. A., Tendler, B. C., & Mars, R. B. (2024). *Comparative neuroimaging*

of the carnivoran brain: Neocortical sulcal anatomy. https://doi.org/10.7554/elife.100851.1
Boelts, J., Deistler, M., Gloeckler, M., Tejero-Cantero, Á., Lueckmann, J.-M., Moss, G., Steinbach, P., Moreau, T., Muratore, F., Linhart, J., & others. (2024). sbi reloaded: a toolkit for simulation-based inference workflows. *ArXiv Preprint ArXiv:2411.17337*.
Bryant, K., Ardesch, D. J., Roumazeilles, L., Scholtens, L. H., Khrapitchev, A. A., Tendler, B. C., Wu, W., Miller, K. L., Sallet, J., den Heuvel, M. P., & others. (2021). Diffusion MRI data, sulcal anatomy, and tractography for eight species from the Primate Brain Bank. *Brain Structure and Function*.
Buxton, R. B. (1993). The diffusion sensitivity of fast steady-state free precession imaging. *Magnetic Resonance in Medicine*. https://doi.org/10.1002/mrm.1910290212
Cardenas, A. M., Sarlls, J. E., Kwan, J. Y., Bageac, D., Gala, Z. S., Danielian, L. E., Ray-Chaudhury, A., Wang, H. W., Miller, K. L., Foxley, S., Jbabdi, S., Welsh, R. C., & Floeter, M. K. (2017). Pathology of callosal damage in ALS: An ex-vivo, 7 T diffusion tensor MRI study. *NeuroImage: Clinical*. https://doi.org/10.1016/j.nicl.2017.04.024
Chenevert, T. L., Brunberg, J. A., & Pipe, J. G. (1990). Anisotropic diffusion in human white matter: demonstration with MR techniques in vivo. *Radiology*, *177*(2), 401–405.
Choe, A. S., Stepniewska, I., Colvin, D. C., Ding, Z., & Anderson, A. W. (2012). Validation of diffusion tensor MRI in the central nervous system using light microscopy: quantitative comparison of fiber properties. *NMR in Biomedicine*, *25*(7), 900–908.
Cook, P. F., & Berns, G. (2022). Volumetric and connectivity assessment of the caudate nucleus in California sea lions and coyotes. *Animal Cognition*, *25*(5), 1231–1240.
Cranmer, K., Brehmer, J., & Louppe, G. (2020). The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, *117*(48), 30055–30062.
Deistler, M., Macke, J. H., & Gonçalves, P. J. (2022). Energy-efficient network activity from disparate circuit parameters. *Proceedings of the National Academy of Sciences*, *119*(44), e2207632119.
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.
Durkan, C., Bekasov, A., Murray, I., & Papamakarios, G. (2019). Neural spline flows. *Advances in Neural Information Processing Systems*, *32*.
Edlow, B. L., Keene, C. D., Perl, D. P., Iacono, D., Folkerth, R. D., Stewart, W., Mac Donald, C. L., Augustinack, J., Diaz-Arrastia, R., Estrada, C., & others. (2018). Multimodal characterization of the late effects of traumatic brain injury: a methodological overview of the late effects of traumatic brain injury project. *Journal of Neurotrauma*, *35*(14), 1604–1619.
Eggl, M. F., & De Santis, S. (2024). More with less: Simulation-based inference enables accurate diffusion-weighted MRI with minimal acquisition time. *BioRxiv*, 2011–2024.
Flem, S., Berns, G., Inglis, B., Niederhut, D., Montie, E., Deacon, T., Miller, K. L., Tyack, P., & Cook, P. F. (2025). Lateralized cerebellar connectivity differentiates auditory pathways in echolocating and non-echolocating whales. *PLoS One*, *20*(6), e0323617.
Foxley, S., Jbabdi, S., Clare, S., Lam, W., Ansorge, O., Douaud, G., & Miller, K. (2014). Improving diffusion-weighted imaging of post-mortem human brains: SSFP at 7T. *NeuroImage*. https://doi.org/10.1016/j.neuroimage.2014.08.014
Freed, D. E., Scheven, U. M., Zielinski, L. J., Sen, P. N., & Hürlimann, M. D. (2001). Steady-state free precession experiments and exact treatment of diffusion in a uniform gradient. *Journal of Chemical Physics*. https://doi.org/10.1063/1.1389859
Greenberg, D., Nonnenmacher, M., & Macke, J. (2019). Automatic posterior transformation for likelihood-free inference. *International Conference on Machine Learning*, 2404–2414.
Gudbjartsson, H., & Patz, S. (1995a). The rician distribution of noisy MRI data. *Magnetic Resonance in Medicine*. https://doi.org/10.1002/mrm.1910340618
Gudbjartsson, H., & Patz, S. (1995b). The rician distribution of noisy MRI data. *Magnetic Resonance in Medicine*, *34*(6), 910–914. https://doi.org/10.1002/mrm.1910340618
Howard, A. F., Mollink, J., Kleinnijenhuis, M., Pallebage-Gamarallage, M., Bastiani, M., Cottaar, M., Miller, K. L., & Jbabdi, S. (2019). Joint modelling of diffusion MRI and microscopy. *NeuroImage*. https://doi.org/10.1016/j.neuroimage.2019.116014
Jallais, M., & Palombo, M. (2024). Introducing $μ$GUIDE for quantitative imaging via generalized uncertainty-driven inference using deep learning. *Elife*, *13*, RP101069.
Jenkinson, M., Bannister, P., Brady, M., & Smith, S. (2002). Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*, *17*(2), 825–841. https://doi.org/10.1016/S1053-8119(02)91132-8

Jenkinson, M., & Smith, S. (2001). A global optimisation method for robust affine registration of brain images. *Medical Image Analysis*, *5*(2), 143–156. https://doi.org/10.1016/S1361-8415(01)00036-6

Ji, H., Wu, Z., Speckert, A., Tuura, R., Hagmann, C., Latal, B., Knirsch, W., & Jakab, A. (2024). Developmental trajectory of thalamus topography during the late preterm and perinatal period in normal development, after premature birth and in congenital heart defects. *BioRxiv*, 2008–2024.

Jones, D. K., & Basser, P. J. (2004). "Squashing peanuts and smashing pumpkins": How noise distorts diffusion-weighted MR data. *Magnetic Resonance in Medicine*. https://doi.org/10.1002/mrm.20283

Jones, E., Oliphant, T., Peterson, P., & Others. (2001). *http://www.SciPy.org*. SciPy: Open Source Scientific Tools for Python.

Kaiser, R., Bartholdi, E., & Ernst, R. R. (1974). Diffusion and field-gradient effects in NMR Fourier spectroscopy. *The Journal of Chemical Physics*. https://doi.org/10.1063/1.1681477

Kellner, E., Dhital, B., Kiselev, V. G., & Reisert, M. (2016). Gibbs-ringing artifact removal based on local subvoxel-shifts. *Magnetic Resonance in Medicine*. https://doi.org/10.1002/mrm.26054

Le Bihan, D. (1988). Intravoxel incoherent motion imaging using steady-state free precession. *Magnetic Resonance in Medicine*. https://doi.org/10.1002/mrm.1910070312

Lueckmann, J.-M., Goncalves, P. J., Bassetto, G., Öcal, K., Nonnenmacher, M., & Macke, J. H. (2017). Flexible statistical inference for mechanistic models of neural dynamics. *Advances in Neural Information Processing Systems*, *30*.

Manzano-Patron, J. P., Deistler, M., Schröder, C., Kypraios, T., Gonçalves, P. J., Macke, J. H., & Sotiropoulos, S. N. (2025). Uncertainty mapping and probabilistic tractography using Simulation-Based Inference in diffusion MRI: A comparison with classical Bayes. *Medical Image Analysis*, 103580.

McNab, J. A., Edlow, B. L., Witzel, T., Huang, S. Y., Bhat, H., Heberlein, K., Feiweier, T., Liu, K., Keil, B., Cohen-Adad, J., & others. (2013). The Human Connectome Project and beyond: initial applications of 300 mT/m gradients. *Neuroimage*, *80*, 234–245.

McNab, J. A., Jbabdi, S., Deoni, S. C. L., Douaud, G., Behrens, T. E. J., & Miller, K. L. (2009). High resolution diffusion-weighted imaging in fixed human brain using diffusion-weighted steady state free precession. *NeuroImage*. https://doi.org/10.1016/j.neuroimage.2009.01.008

McNab, J. A., & Miller, K. L. (2008). Sensitivity of diffusion weighted steady state free precession to anisotropic diffusion. *Magnetic Resonance in Medicine*. https://doi.org/10.1002/mrm.21668

McNab, J. A., & Miller, K. L. (2010). Steady-state diffusion-weighted imaging: Theory, acquisition and analysis. In *NMR in Biomedicine*. https://doi.org/10.1002/nbm.1509

Merboldt, K.-D., Hänicke, W., Gyngell, M. L., Frahm, J., & Bruhn, H. (1989a). Rapid NMR imaging of molecular self-diffusion using a modified CE-FAST sequence. *Journal of Magnetic Resonance (1969)*. https://doi.org/10.1016/0022-2364(89)90170-4

Merboldt, K.-D., Hänicke, W., Gyngell, M. L., Frahm, J., & Bruhn, H. (1989b). The influence of flow and motion in MRI of diffusion using a modified CE-FAST sequence. *Magnetic Resonance in Medicine*, *12*(2), 198–208.

Miller, K. L., Hargreaves, B. A., Gold, G. E., & Pauly, J. M. (2004). Steady-State Diffusion-Weighted Imaging of In Vivo Knee Cartilage. *Magnetic Resonance in Medicine*. https://doi.org/10.1002/mrm.10696

Miller, K. L., McNab, J. A., Jbabdi, S., & Douaud, G. (2012). Diffusion tractography of post-mortem human brains: Optimization and comparison of spin echo and steady-state free precession techniques. *NeuroImage*. https://doi.org/10.1016/j.neuroimage.2011.09.054

Mollink, J., Kleinnijenhuis, M., Cappellen van Walsum, A. M. van, Sotiropoulos, S. N., Cottaar, M., Mirfin, C., Heinrich, M. P., Jenkinson, M., Pallebage-Gamarallage, M., Ansorge, O., Jbabdi, S., & Miller, K. L. (2017). Evaluating fibre orientation dispersion in white matter: Comparison of diffusion MRI, histology and polarized light imaging. *NeuroImage*. https://doi.org/10.1016/j.neuroimage.2017.06.001

Moseley, M. E., Cohen, Y., Kucharczyk, J., Mintorovitch, J., Asgari, H. S., Wendland, M. F., Tsuruda, J., & Norman, D. (1990). Diffusion-weighted MR imaging of anisotropic water diffusion in cat central nervous system. *Radiology*, *176*(2), 439–445.

Nolan, A. L., Petersen, C., Iacono, D., Mac Donald, C. L., Mukherjee, P., van der Kouwe, A., Jain, S., Stevens, A., Diamond, B. R., Wang, R., & others. (2021). Tractography-pathology correlations in traumatic brain injury: a TRACK-TBI study. *Journal of Neurotrauma*, *38*(12), 1620–1631.

Novikov, D. S., Fieremans, E., Jespersen, S. N., & Kiselev, V. G. (2019). Quantifying brain microstructure with diffusion MRI: Theory and parameter estimation. *NMR in Biomedicine*, *32*(4), e3998.

O'Halloran, R., Aksoy, M., Aboussouan, E., Peterson, E., Van, A., & Bammer, R. (2015). Real-time correction of rigid body motion-induced phase errors for diffusion-weighted steady-state free precession imaging. *Magnetic Resonance in Medicine*, *73*(2), 565–576.

O'Halloran, R. L., Aksoy, M., Van, A. T., & Bammer, R. (2013). 3D isotropic high-resolution diffusion-weighted

MRI of the whole brain with a motion-corrected steady-state free precession sequence. *Magnetic Resonance in Medicine*, *70*(2), 466–478.

Orekhova, K., Selmanovic, E., De Gasperi, R., Gama Sosa, M. A., Wicinski, B., Maloney, B., Seifert, A., Alipour, A., Balchandani, P., Gerussi, T., & others. (2022). Multimodal assessment of bottlenose dolphin auditory nuclei using 7-Tesla MRI, immunohistochemistry and stereology. *Veterinary Sciences*, *9*(12), 692.

Papamakarios, G., & Murray, I. (2016). Fast ε-free inference of simulation models with bayesian conditional density estimation. *Advances in Neural Information Processing Systems*, *29*.

Papamakarios, G., Nalisnick, E., Rezende, D. J., Mohamed, S., & Lakshminarayanan, B. (2021). Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, *22*(57), 1–64.

Radev, S. T., Schmitt, M., Schumacher, L., Elsemüller, L., Pratz, V., Schälte, Y., Köthe, U., & Bürkner, P.-C. (2023). BayesFlow: Amortized Bayesian workflows with neural networks. *ArXiv Preprint ArXiv:2306.16015*.

Roebroeck, A., Miller, K. L., & Aggarwal, M. (2019). Ex vivo diffusion MRI of the human brain: Technical challenges and recent advances. *NMR in Biomedicine*. https://doi.org/10.1002/nbm.3941

Roumazeilles, L., Eichert, N., Bryant, K. L., Folloni, D., Sallet, J., Vijayakumar, S., Foxley, S., Tendler, B. C., Jbabdi, S., Reveley, C., Verhagen, L., Dershowitz, L. B., Guthrie, M., Flach, E., Miller, K. L., & Mars, R. B. (2020). Longitudinal connections and the organization of the temporal cortex in macaques, great apes, and humans. *PLoS Biology*. https://doi.org/10.1371/journal.pbio.3000810

Schilling, K. G., Howard, A. F. D., Grussu, F., Ianus, A., Hansen, B., Barrett, R. L. C., Aggarwal, M., Michielse, S., Nasrallah, F., Syeda, W., & others. (2025). Considerations and recommendations from the ISMRM Diffusion Study Group for preclinical diffusion MRI: Part 3—Ex vivo imaging: Data processing, comparisons with microscopy, and tractography. *Magnetic Resonance in Medicine*, *93*(6), 2561–2582.

Stepišnik, J. (1993). Time-dependent self-diffusion by NMR spin-echo. *Physica B: Condensed Matter*, *183*(4), 343–350.

Szafer, A., Zhong, J., Anderson, A. W., & Gore, J. C. (1995). Diffusion-weighted imaging in tissues: theoretical models. *NMR in Biomedicine*, *8*(7), 289–296.

Tejero-Cantero, A., Boelts, J., Deistler, M., Lueckmann, J.-M., Durkan, C., Gonçalves, P. J., Greenberg, D. S., & Macke, J. H. (2020). sbi: A toolkit for simulation-based inference. *Journal of Open Source Software*, *5*(52), 2505. https://doi.org/10.21105/joss.02505

Tendler, B. C. (2025). Investigating time-independent and time-dependent diffusion phenomena using steady-state diffusion MRI. *Scientific Reports*, *15*(1), 3580.

Tendler, B. C. (2023). Theoretical framework to characterise motion artefacts arising from steady-state diffusion MRI. *European Society of Magnetic Resonance in Medicine & Biology (ESMRMB) Annual Scientific Meeting*, LB298.

Tendler, B. C., Foxley, S., Cottaar, M., Jbabdi, S., & Miller, K. L. (2020). Modeling an equivalent b-value in diffusion-weighted steady-state free precession. *Magnetic Resonance in Medicine*. https://doi.org/10.1002/mrm.28169

Tendler, B. C., Foxley, S., Hernandez-Fernandez, M., Cottaar, M., Scott, C., Ansorge, O., Miller, K. L., & Jbabdi, S. (2020). Use of multi-flip angle measurements to account for transmit inhomogeneity and non-Gaussian diffusion in DW-SSFP. *NeuroImage*. https://doi.org/10.1016/j.neuroimage.2020.117113

Tendler, B. C., Hanayik, T., Ansorge, O., Bangerter-Christensen, S., Berns, G. S., Bertelsen, M. F., Bryant, K. L., Foxley, S., van den Heuvel, M. P., Howard, A. F. D., & others. (2022). The Digital Brain Bank, an open access platform for post-mortem imaging datasets. *Elife*, *11*, e73153.

Tendler, B. C., Qi, F., Foxley, S., Pallebage-Gamarallage, M., Menke, R. A. L., Ansorge, O., Hurley, S. A., & Miller, K. L. (2021). A method to remove the influence of fixative concentration on postmortem T2 maps using a kinetic tensor model. *Human Brain Mapping*, *42*(18), 5956–5972. https://doi.org/https://doi.org/10.1002/hbm.25661

Tendler, B. C., Warrington, S., Selim, M. K., Wu, W., Adriany, G., Auerbach, E. J., Bratch, A., Farooq, H., Harel, N., Heilbronner, S., Jbabdi, S., Jungst, S., Lenglet, C., Moeller, S., Pestilli, F., Pisharady, P., Ugurbil, K., Waks, M., Yacoub, E., … Zimmermann, J. (2025). Diffusion MRI Acquisition Methods For Post-Mortem Imaging On A 10.5 T Whole-Body Human MRI Scanner. *Proceedings of the 33rd Annual Meeting of ISMRM*, 4069.

van Veluw, S. J., Reijmer, Y. D., van der Kouwe, A. J., Charidimou, A., Riley, G. A., Leemans, A., Bacskai, B. J., Frosch, M. P., Viswanathan, A., & Greenberg, S. M. (2019). Histopathology of diffusion imaging abnormalities in cerebral amyloid angiopathy. *Neurology*, *92*(9), e933--e943.

Vangelderen, P., DesPres, D., Vanzijl, P. C. M., & Moonen, C. T. W. (1994). Evaluation of restricted diffusion in cylinders. Phosphocreatine in rabbit leg muscle. *Journal of Magnetic Resonance, Series B*, *103*(3), 255–260.

Verma, Y., Bharti, A., & Garg, V. (2025). Robust simulation-based inference under missing data via neural processes. *ArXiv Preprint ArXiv:2503.01287*.

Warrington, S., Tendler, B. C., Selim, M. K., Wu, W., Farooq, H., Moeller, S., Pisharady, P., Auerbach, E. J., Adriany, G., Bratch, A., Waks, M., Jungst, S., Harel, N., Pestilli, F., Heilbronner, S., Yacoub, E., Ugurbil, K., Lenglet, C., Miller, K., … Sotiropoulos, S. (2025). High Resolution Diffusion MRI Tractography In The In-Vivo & Ex-Vivo NHP Brain Using A Human 10.5T Scanner. *Proceedings of the 33rd Annual Meeting of ISMRM*, 0322.

Weigel, M. (2015). Extended phase graphs: Dephasing, RF pulses, and echoes - Pure and simple. In *Journal of Magnetic Resonance Imaging*. https://doi.org/10.1002/jmri.24619

Weigel, M., Schwenk, S., Kiselev, V. G., Scheffler, K., & Hennig, J. (2010). Extended phase graphs with anisotropic diffusion. *Journal of Magnetic Resonance*, *205*(2), 276–285.

Wilkinson, M., Wang, R., van der Kouwe, A., & Takahashi, E. (2016). White and gray matter fiber pathways in autism spectrum disorder revealed by ex vivo diffusion MR tractography. *Brain and Behavior*. https://doi.org/10.1002/brb3.483

Yarnykh, V. L. (2007). Actual flip-angle imaging in the pulsed steady state: A method for rapid three-dimensional mapping of the transmitted radiofrequency field. *Magnetic Resonance in Medicine*. https://doi.org/10.1002/mrm.21120

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

Zhang, Z. W., Song, L. J., Meng, Q. F., Li, Z. P., Luo, B. N., Yang, Y. H., & Pei, Z. (2008). High-resolution diffusion-weighted MR imaging of the human lumbosacral plexus and its branches based on a steady-state free precession imaging technique at 3T. *American Journal of Neuroradiology*, *29*(6), 1092–1094.

Zheng, Z., Tachrount, M., Miller, K. L., Cottaar, M., & Tendler, B. C. (2024). A modelling and experimental framework to investigate the sensitivity of steady-state diffusion MRI to microstructure. *Proceedings of the 32nd Annual Meeting of ISMRM*, 0584.

Zhu, Q., Li, S., Cao, Z., Shen, Y., Xu, H., Xu, G., Li, H., Zhu, K., Zhao, Z., Zhang, J., & others. (2024). 7 Tesla multimodal MRI dataset of ex-vivo human brain. *ArXiv Preprint ArXiv:2412.17816*.

Zhu, S., Huszar, I. N., Cottaar, M., Daubney, G., Eichert, N., Hanayik, T., Khrapitchev, A. A., Mars, R. B., Mollink, J., Sallet, J., & others. (2025). Imaging the structural connectome with hybrid MRI-microscopy tractography. *Medical Image Analysis*, 103498.

Appendix 1

Based on the proposed analytical model in (Freed et al., 2001), the DW-SSFP sequence incorporating a diffusion gradient of duration δ is modelled as:

$$S_{\mathrm{DW-SSFP}} = -r_1 \cdot b_0, \qquad \text{[A1]}$$

where:

$$b_0 = \frac{-\sin\alpha \cdot S_0(1 - E_{1,0})}{A_{p=0} - B_{p=0} + E_{2,-1} \cdot C_{p=0} \cdot r_1} \qquad \text{[A2]}$$

$$r_1 = \frac{1}{E_{2,-1} \cdot B_{p=0}} x_1 + \frac{E_{2,0} \cdot C_{p=1}}{B_{p=1}} \qquad \text{[A3]}$$

defining:

$$A_p = \frac{1}{2} \cdot \left(E_{1,p} - 1\right) \cdot (1 + \cos\alpha) \qquad \text{[A4]}$$

$$B_p = \frac{1}{2} \cdot \left(E_{1,p} + 1\right) \cdot (1 - \cos\alpha) \qquad \text{[A5]}$$

$$C_p = E_{1,p} - \cos\alpha \qquad \text{[A6]}$$

$$E_{1,p} = e^{-\frac{\mathrm{TR}}{T_1} - D \cdot \gamma^2 \cdot G^2 \cdot \delta^2 \cdot p^2 \cdot \mathrm{TR}} \qquad \text{[A7]}$$

$$E_{2,p} = e^{-\frac{\mathrm{TR}}{T_2} - D \cdot \gamma^2 \cdot G^2 \cdot \delta^2 \cdot \left(\left(\mathrm{p}^2 + \mathrm{p} + \frac{1}{3}\right) \cdot \delta + (\mathrm{p}+1)^2 \cdot (\mathrm{TR} - \delta)\right)} \qquad \text{[A8]}$$

$$x_1 = \frac{n_1}{d_1 + \frac{n_2}{d_2 + \frac{n_3}{d_3 + \cdots + \frac{n_l}{d_l + e_l}}}}, \qquad \text{[A9]}$$

with:

$$n_p = -E_{2,-p} \cdot E_{2,p-1} \cdot A_p^2 \cdot \frac{B_{p-1}}{B_p} \qquad \text{[A10]}$$

$$d_p = \left(A_p - B_p\right) + E_{2,-p-1} \cdot E_{2,p} \cdot \frac{B_p}{B_{p+1}} \cdot C_{p+1} \qquad \text{[A11]}$$

$$e_p = -E_{2,-p-1} \cdot E_{2,p} \cdot \frac{B_p}{B_{p+1}} \cdot C_{p+1}. \qquad \text{[A12]}$$

Transmit (B_1) inhomogeneity can be incorporated as a modulation of the nominal flip angle α.

Appendix 2

The DW-SSFP analytical model with integrated diffusion tensor is equivalent to Appendix 1, setting:

$$D = \hat{g}_k^T \boldsymbol{D} \hat{g}_k, \qquad \text{[A13]}$$

where $\hat{g}_k$ is the gradient orientation and $\boldsymbol{D}$ is the diffusion tensor.