



HYDI-DSI revisited: Constrained non-parametric EAP imaging without q-space re-gridding

Antonio Tristán-Vega^{a,*}, Tomasz Pieciak^{a,b}, Guillem París^a, Justino R. Rodríguez-Galván^a, Santiago Aja-Fernández^a

^a LPI, ETSI Telecomunicación, Universidad de Valladolid, Spain

^b AGH University of Science and Technology, Kraków, Poland

ARTICLE INFO

Keywords:

Diffusion MRI
Hybrid diffusion imaging
Diffusion spectrum imaging
Ensemble average propagator

ABSTRACT

Hybrid Diffusion Imaging (HYDI) was one of the first attempts to use multi-shell samplings of the q-space to infer diffusion properties beyond Diffusion Tensor Imaging (DTI) or High Angular Resolution Diffusion Imaging (HARDI). HYDI was intended as a flexible protocol embedding both DTI (for lower b -values) and HARDI (for higher b -values) processing, as well as Diffusion Spectrum Imaging (DSI) when the entire data set was exploited. In the latter case, the spherical sampling of the q-space is re-gridded by interpolation to a Cartesian lattice whose extent covers the range of acquired b -values, hence being acquisition-dependent. The Discrete Fourier Transform (DFT) is afterwards used to compute the corresponding Cartesian sampling of the Ensemble Average Propagator (EAP) in an entirely non-parametric way. From this lattice, diffusion markers such as the Return To Origin Probability (RTOP) or the Mean Squared Displacement (MSD) can be numerically estimated.

We aim at re-formulating this scheme by means of a Fourier Transform encoding matrix that eliminates the need for q-space re-gridding at the same time it preserves the non-parametric nature of HYDI-DSI. The encoding matrix is adaptively designed at each voxel according to the underlying DTI approximation, so that an optimal sampling of the EAP can be pursued without being conditioned by the particular acquisition protocol. The estimation of the EAP is afterwards carried out as a regularized Quadratic Programming (QP) problem, which allows to impose positivity constraints that cannot be trivially embedded within the conventional HYDI-DSI. We demonstrate that the definition of the encoding matrix in the adaptive space allows to analytically (as opposed to numerically) compute several popular descriptors of diffusion with the unique source of error being the cropping of high frequency harmonics in the Fourier analysis of the attenuation signal. They include not only RTOP and MSD, but also Return to Axis/Plane Probabilities (RTAP/RTPP), which are defined in terms of specific spatial directions and are not available with the former HYDI-DSI. We report extensive experiments that suggest the benefits of our proposal in terms of accuracy, robustness and computational efficiency, especially when only standard, non-dedicated q-space samplings are available.

1. Introduction

Diffusion Magnetic Resonance Imaging (dMRI) allows probing the random movement of water molecules inside human tissues, especially the white matter of the brain, *in vivo*. The basic dMRI sequence comprises two pulsed gradients with duration δ taken apart an idle time Δ , with a re-focusing 180° RF pulse in between (Stejskal and Tanner, 1965). This way, water molecules moving along the applied gradient will experiment a net de-phasing that translates in the attenuation of the acquired T2 signal with respect to the unweighted one (the so-called baseline). The *strength* of this effect, characterized by the so-called b -value, increases linearly with Δ , and quadratically with δ and the magnitude $\|G\|$ of the gradient applied.

This kind of contrast has been thoroughly used in clinical research to characterize a number of pathologies such as Alzheimer's disease (Dai et al., 2015; Fick et al., 2016a), Parkinson's disease (Aja-Fernández et al., 2020), Traumatic Brain Injury (TBI) (Wu et al., 2018; Muller et al., 2021), stroke (Boscolo Galazzo et al., 2018; Brusini et al., 2016), multiple sclerosis (Hosseinbor et al., 2012), migraine (Planchuelo-Gómez et al., 2020, 2021), and many others. Though a minimum of six gradient directions with a constant b -value suffices to probe the meso-structural anisotropy of the white matter, the advent of more sophisticated MRI machinery and acquisition protocols has led into the routine acquisition of advanced data sets with several hundreds of

* Corresponding author.

E-mail address: atriveg@lpi.tel.uva.es (A. Tristán-Vega).

<https://doi.org/10.1016/j.media.2022.102728>

Received 26 January 2022; Received in revised form 20 October 2022; Accepted 7 December 2022

Available online 9 December 2022

1361-8415/© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

diffusion gradients with varying directions but also varying b-values. Among them, multi-shell data sets, i.e. acquisitions where the acquired gradients are arranged in a regular spherical lattice, have rapidly become the standard. Indeed, several databases with both healthy and diseased subjects have been publicly issued in the last few years (Jack et al., 2008; Van Essen et al., 2013; Ziegler et al., 2014; Fan et al., 2016; Froeling et al., 2017; Tax et al., 2019; Koller et al., 2021).

Two complementary approaches have focused the recent research on advanced dMRI (Novikov et al., 2018): the first one aims at modeling the diffusion signal as a mixture of micro-structural compartments whose responses can be individually modeled, so that their partial volume fractions and individual features can be disentangled from the diffusion measurements. This group includes, to name but a few, the Composite Hindered and Restricted Model for Diffusion (Assaf and Basser, 2005, CHARMED), the Neurite Orientation Dispersion and Density Imaging (Zhang et al., 2012, NODDI), or the Multi-tissue Constrained Spherical Deconvolution (Jeurissen et al., 2014, CSD). These methods have the advantage of providing directly interpretable micro-structural features, and also explaining both the low and the high b-value regimen of the diffusion signal and, in particular, its slow decaying tails. On the other hand, they give rise to ill-posed optimization problems that are often simplified resorting to modeling assumptions such as minimum tortuosity (Zhang et al., 2012; Kaden et al., 2016). These suppositions, however, have been empirically evidenced to largely deviate from reality (Lampinen et al., 2020).

The second group of techniques generalize the *classical* Diffusion Tensor Imaging (DTI) approach (Basser et al., 1994; Basser and Pierpaoli, 1996) by drawing signal representations, as opposed to geometrical models, such as mono-exponentials (Aja-Fernández et al., 2021, 2022), multi-exponentials (Pfeuffer et al., 1999; Benjamini and Basser, 2019), or DTI distributions (Topgaard, 2017). In particular, Ensemble Average Propagator (EAP) imaging is a very active research field within this second trend: the EAP, $P(\mathbf{R})$, is a positive, unit-mass, antipodal-symmetric probability density function related to the positive, antipodal-symmetric attenuation signal $E(\mathbf{q}) = S(\mathbf{q})/S_0$ (with $S(\mathbf{q})$ the signal acquired when a gradient wave-vector \mathbf{q} , $\|\mathbf{q}\| \propto \delta \|\mathbf{G}\|$, is applied, and S_0 the unweighted T2 baseline image), as a pair of Fourier transforms (Callaghan, 1991):

$$\begin{aligned} P(\mathbf{R}) &= \iiint_{\mathbb{R}^3} E(\mathbf{q}) \exp(-j2\pi\mathbf{q}^T\mathbf{R}) d\mathbf{q} \\ \longleftrightarrow E(\mathbf{q}) &= \iiint_{\mathbb{R}^3} P(\mathbf{R}) \exp(j2\pi\mathbf{q}^T\mathbf{R}) d\mathbf{R}, \end{aligned} \quad (1)$$

where $\tau = \Delta - \delta/3$ is the effective diffusion time. For the previous equation to be fulfilled in terms of classical functions, $E(\mathbf{q})$ must rapidly vanish as $\|\mathbf{q}\| \rightarrow \infty$, so that it remains square-integrable and $P(\mathbf{R})$ does not comprise Dirac's delta-like singularities (Oppenheim et al., 1997). As pointed out by Novikov et al. (2018), this issue makes EAP imaging inherently incompatible with actual models for intra-axonal diffusion, as long as the signal produced by confined water compartments is heavy-tailed. Within DTI, the EAP is represented as a zero-mean Gaussian process with covariance matrix D , which is a 3×3 rank-2 tensor, symmetric, positive (semi)-definite, and independent on the effective diffusion time τ (Basser et al., 1994; Basser and Pierpaoli, 1996):

$$\begin{aligned} P_G(\mathbf{R}) &= \frac{1}{\sqrt{\det(D)(4\pi\tau)^3}} \exp\left(-\frac{\mathbf{R}^T D^{-1} \mathbf{R}}{4\tau}\right) \\ \longleftrightarrow E_G(\mathbf{q}) &= \exp(-4\pi^2\tau\mathbf{q}^T D \mathbf{q}), \end{aligned} \quad (2)$$

where the b-value is defined as $b = 4\pi^2\tau\|\mathbf{q}\|^2$. Diffusion Spectrum Imaging (DSI) was aimed at avoiding the limitations of the model in Eq. (2) by attaining a non-parametric sampling of $P(\mathbf{R})$ at a regular Cartesian lattice (Wedeen et al., 2005). The straightforward approach is sampling $E(\mathbf{q})$ itself in a regular lattice and using the Discrete Fourier

Transform (DFT) to find estimates of $P(\mathbf{R})$, which implies acquiring a huge amount of q-space samples to avoid aliasing artifacts. Alternatively, Compressed Sensing (CS) techniques can be used to vastly reduce this demand (Bilgic et al., 2012; Menzel et al., 2011; Young et al., 2017) as long as a *non-coherent* sampling of the whole q-space is available and the EAP can be sparsely represented in a certain function basis (Donoho, 2006). Since we are mostly interested in multi-shell samplings with a relatively small maximum b-value, the former requirement is hardly met.

Many EAP imaging techniques circumvent the lack of a detailed q-space sampling by assuming the EAP and the attenuation signal can be faithfully represented as a superposition of pre-designed basis functions (dictionary atoms) whose 3-D Fourier transforms may be easily characterized. This approach includes multiple q-shell Diffusion Propagator Imaging (Descoteaux et al., 2011, mq-DPI), Bessel-Fourier Orientation Reconstruction (Hosseinbor et al., 2013, BFOR), Spherical Polar Fourier reconstruction (Assemlal et al., 2009; Merlet and Deriche, 2013, SPF), Simple Harmonic Oscillator based Reconstruction and Estimation (Özarslan et al., 2013a, SHORE), Mean Apparent Propagator MRI (Özarslan et al., 2013b, MAP-MRI) and its Laplacian-regularized version (Fick et al., 2016b, MAPL), or directional Radial Basis Functions (Ning et al., 2015, RBF). Patch-based dictionary learning, as opposed to dictionary design, has also been used to low-rank represent the attenuation signal (Vemuri et al., 2019).

The general idea behind these techniques is that the attenuation signal $E(\mathbf{q})$ can be written as the linear superposition of a relatively small number of parametric continuous functions, whose shape and size parameters can be either pre-defined (Assemlal et al., 2009; Merlet and Deriche, 2013), learned (Vemuri et al., 2019), adaptively fitted depending on the diffusion profile at each voxel (Özarslan et al., 2013b; Fick et al., 2016b), or even dynamically computed at the same time as the coefficients of the linear combination (Ning et al., 2015). Afterwards, the linearity of the Fourier transform can be exploited to apply the same linear combination to the Fourier transforms of the basis functions and conversely represent the EAP as a linear mixture of parametric continuous functions. Provided the reduced number of q-space samples, together with the poor Signal-to-Noise Ratio (SNR) commonly found in dMRI volumes, it is common to find regularization penalties in the estimation of the linear coefficients that describe $E(\mathbf{q})$, as well as constraints such as positivity or unit mass of the EAP (Dela Haije et al., 2020). More recently, the Micro-Structure adaptive kernels and dual Fourier Integral Transforms (Tristán-Vega and Aja-Fernández, 2021, MiSFIT) has allowed to dramatically reduce the complexity of EAP imaging by representing it as the spherical convolution of a fiber Orientation Distribution Function (fODF) with a Gaussian kernel. Nonetheless, this technique applies only to multi-shell acquisitions, but not to more general protocols.

All the above methods share the same philosophy of representing the low-pass EAP parametrically, as a mixture of continuous domain functions. As opposed, DSI and Hybrid Diffusion Imaging-based DSI (Wu et al., 2008; Wu and Alexander, 2007, HYDI-DSI) tackle the problem in a different, straightforward way, as it is depicted in Fig. 1 (top): the (still low-pass) EAP is no longer represented by means of a collection of parameters describing a continuous mixture, but instead a discrete sampling of the EAP at a regular 3-D lattice (typically $9 \times 9 \times 9$) is pursued. To that end, a corresponding lattice is defined over the attenuation signal domain (ii), whose support comes determined by the maximum b-value, b_{\max} , acquired in the multi-shell sampling (i), i.e. $E(\mathbf{q})$ is assumed to live inside $\Omega = [-q_{\max}, q_{\max}]^3 \subset \mathbb{R}^3$, where $q_{\max} = \sqrt{b_{\max}/4\pi^2\tau}$. The values of $E(\mathbf{q})$ at these lattice points are obtained by means of grid interpolation: the convex hull of the sampled data is calculated (Barber et al., 1996), and Delaunay triangulation is used to parcel it (Watson, 1994). The problem then reduces to the linear interpolation of each lattice node depending on the parcel it lies within. Finally, the DFT can be computed to retrieve a $9 \times 9 \times 9$ lattice in the

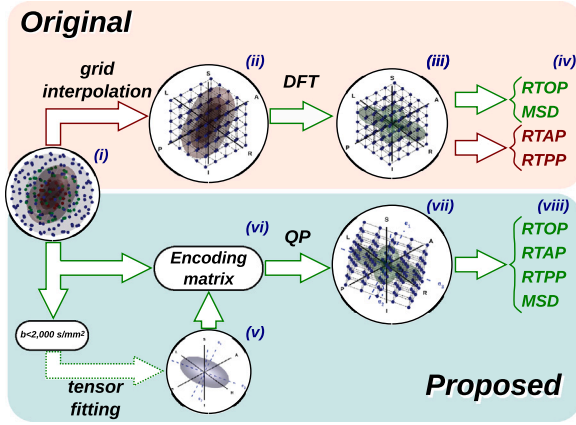


Fig. 1. The original HYDI-DSI vs. our proposal. Red arrows highlight those steps in the original method resorting to numerical interpolation (those we aim to avoid).

EAP domain (iii) with spatial resolution $(1/2q_{\max})^3$ (Oppenheim et al., 1997).

Despite its success in describing the anatomy of the white matter in many situations such as TBI (Wu et al., 2018; Muller et al., 2021), Alzheimer's disease (Daianu et al., 2015), or gender/age-related changes (Wu et al., 2011), this scheme is not free of certain problems. The first one is that the bandwidth of the EAP, defined by Ω , directly depends on the acquisition protocol, and more specifically on the maximum b-value acquired. Indeed, the original HYDI-DSI imposes a specific acquisition protocol with evenly spaced shells up to $b_{\max} \approx 10,000$ s/mm², so that the lattice interpolation is rather uniform (meaning that each measured q-space sample will be used equally often for interpolation). Additionally, though the unit-mass constraint of the EAP is guaranteed by the DFT operator by simply placing $E(\mathbf{0}) = 0$, positivity constraints cannot be pursued. Such constraints have been proven especially meaningful within dMRI (Dela Haije et al., 2020). Finally, the analytical computations of the popular Return to Origin Probability (RTOP) or Mean Squared Displacement (MSD) are straightforward with HYDI-DSI, but the Return to Axis/Plane Probabilities (RTAP/RTPP) need to be computed as line/plane integrals in the EAP domain, leading to the need for interpolation of the computed lattice values (iv).

In the present paper, we explore an alternative approach to get rid of q-space interpolation while maintaining the non-parametric nature of HYDI-DSI, see Fig. 1 (bottom). Following the same idea by Özarslan et al. (2013b), the DTI representation is used to describe the low b-value regime of the attenuation signal (v): the eigenvectors and eigenvalues at each voxel are used to rotate and stretch the 3-D space so that the EAP is estimated in an adaptive grid determined by the diffusion properties of the voxel instead of the gradients table (vii). The aim is palliating the dependency of the estimated EAP with the acquisition protocol. To avoid the need for q-space interpolation, the sparsely sampled q-space (i) is analytically related to the adaptive EAP lattice (vii) by means of the so-called **encoding matrix** (vi). This way, the nodal values of the EAP can be directly solved from the q-space samples by solving a Quadratic Programming (QP) problem. The advantage of doing so is two-fold: first, not only unit mass, but also positivity constraints, can be imposed to the EAP. Additionally, by avoiding the computation of the convex hull and Delaunay's triangulation we attain a vast improvement of the overall computational efficiency. Besides, we derive a regularization term based on the energy of the Laplacian of the EAP to deal with low SNR and/or poor sampling density of the q-space. Finally, we demonstrate that the adaptive grid strategy allows the analytical computation of all the RTxP and MSD indices without any further interpolation (viii).

The remainder of the paper is organized as follows. Section 2 details steps v–vii in Fig. 1 (bottom), namely: the arrangement of the

adaptive 3-D lattice based on the DTI approximation (Section 2.1), the definition of the encoding matrix (Sections 2.2 and 2.3), and the QP problem statement (Section 2.4). Section 3 details step viii, i. e. the computation of the RTxP and MSD. Section 4 describes the numerical implementation of the proposed method and the parameters involved. The qualitative and quantitative evaluation of the proposal is addressed in Section 5. Finally, in Section 6 we provide some additional insights into the potential and the limitations of our proposal, as well as its differences and similarities with the related state of the art.

2. EAP reconstruction from scattered multi-shell data

2.1. Adapting the Cartesian grid to the voxel properties

The DTI model provides a good approximation of the diffusion process for b-values under 2000 s/mm². We will assume our q-space sampling includes at least one shell suitable for DTI, so that a diffusion tensor D can be estimated at each voxel to accurately describe the low b-values regime:

$$D = \Theta \Lambda \Theta^T; \quad \Theta = [\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3]; \quad \Lambda = \begin{bmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{bmatrix}, \quad (3)$$

where $0 \leq \lambda_1 \leq \lambda_2 \leq \lambda_3$ are the three real, non-negative eigenvalues of D and $\mathbf{u}_1, \mathbf{u}_2$, and \mathbf{u}_3 are their respective 3×1 , unit-norm, mutually orthogonal eigenvectors. Besides, we will force \mathbf{u}_3 to have the proper orientation so that $\mathbf{u}_3 = \mathbf{u}_1 \times \mathbf{u}_2$ and Θ is a rotation matrix (hence $\det(\Theta) = 1$). We can use this rotation afterwards to align the Cartesian lattice where the EAP will be sampled with the principal directions of D : with the change of variable: $\mathbf{R}' = \Theta^T \mathbf{R} \Leftrightarrow \mathbf{R} = \Theta \mathbf{R}'$, the maximum diffusion direction will become aligned with the 'z' axis, while the secondary diffusion directions will align with 'x' and 'y':

$$\begin{aligned} P_{\Theta}(\mathbf{R}') &= P(\Theta \mathbf{R}') \Leftrightarrow P(\mathbf{R}) = P_{\Theta}(\Theta^T \mathbf{R}); \\ E_{\Theta}(\mathbf{q}') &= \iiint_{\mathbb{R}^3} P_{\Theta}(\mathbf{R}') \exp(j2\pi \mathbf{q}'^T \mathbf{R}') d\mathbf{R}' \\ &\stackrel{\mathbf{R}' = \Theta^T \mathbf{R}}{=} \iiint_{\mathbb{R}^3} P(\mathbf{R}) \exp(j2\pi \mathbf{q}'^T \Theta^T \mathbf{R}) d\mathbf{R} = E(\Theta \mathbf{q}'). \end{aligned} \quad (4)$$

The meaning of the previous equation is that, without any loss of generality, we can assume the Gaussian approximation of the EAP is aligned with the Cartesian axes ('z' being the maximum diffusion direction). It suffices to apply a voxel-dependent rotation Θ^T to the 'gradients table' of the multi-shell sampling, $\mathbf{u}' = \Theta^T \mathbf{u}$, so that:

$$\begin{aligned} P_{G,\Theta}(\mathbf{R}') &= \frac{1}{\sqrt{\det(\Lambda)} (4\pi\tau)^3} \exp\left(-\frac{\mathbf{R}'^T \Lambda^{-1} \mathbf{R}'}{4\tau}\right) \\ &= \frac{1}{\sqrt{\lambda_1 \lambda_2 \lambda_3} (4\pi\tau)^3} \exp\left(\frac{-x^2}{4\tau\lambda_1}\right) \exp\left(\frac{-y^2}{4\tau\lambda_2}\right) \exp\left(\frac{-z^2}{4\tau\lambda_3}\right); \\ E_{G,\Theta}(\mathbf{q}') &= \exp\left(-4\pi^2 \tau \mathbf{q}'^T \Lambda \mathbf{q}'\right) \\ &= \exp\left(-4\pi^2 \tau \lambda_1 q_x^2\right) \exp\left(-4\pi^2 \tau \lambda_2 q_y^2\right) \exp\left(-4\pi^2 \tau \lambda_3 q_z^2\right), \end{aligned} \quad (5)$$

for $\mathbf{R}' = [x, y, z]^T$ and $\mathbf{q}' = [q_x, q_y, q_z]^T$.

2.2. Relating the (rotated) q-space to the (rotated) EAP domain

We will assume the attenuation signal is compact supported, i. e. the value of $E_{\Theta}(\mathbf{q}')$ vanishes to zero outside the 3-D domain $\Omega = (-\frac{Q_x}{2}, \frac{Q_x}{2}) \times (-\frac{Q_y}{2}, \frac{Q_y}{2}) \times (-\frac{Q_z}{2}, \frac{Q_z}{2}) \subset \mathbb{R}^3$. This allows to arrange a 3-D tiling with shifted versions of $E_{\Theta}(\mathbf{q}')$ to build a periodic signal in the three coordinates $\{q_x, q_y, q_z\}$:

$$\begin{aligned} \tilde{q}_u &= q_u - Q_u \left\lfloor \frac{q_u - Q_u/2}{Q_u} \right\rfloor \quad \text{for } u \in \{x, y, z\} \\ \Rightarrow \tilde{E}_{\Theta}(\mathbf{q}') &= E_{\Theta}(\tilde{\mathbf{q}}'), \quad \text{for } \tilde{\mathbf{q}}' = [\tilde{q}_x, \tilde{q}_y, \tilde{q}_z]^T. \end{aligned} \quad (6)$$

Assuming \tilde{E}_θ has finite power, it can be written in terms of its 3-D Fourier series expansion (Oppenheim et al., 1997):

$$\tilde{E}_\theta(\mathbf{q}') = \sum_{k=-\infty}^{\infty} \sum_{l=-\infty}^{\infty} \sum_{m=-\infty}^{\infty} c_{k,l,m} \exp\left(j2\pi\left(\frac{k}{Q_x}q_x + \frac{l}{Q_y}q_y + \frac{m}{Q_z}q_z\right)\right), \quad (7)$$

where the coefficients $c_{k,l,m}$ are computed by projecting \tilde{E}_θ onto each normalized basis function:

$$\begin{aligned} c_{k,l,m} &= \frac{1}{Q} \iiint_{\Omega} \tilde{E}_\theta(\mathbf{q}') \exp\left(-j2\pi\left(\frac{k}{Q_x}q_x + \frac{l}{Q_y}q_y + \frac{m}{Q_z}q_z\right)\right) d\mathbf{q}' \\ &= \frac{1}{Q} \iiint_{\mathbb{R}^3} E_\theta(\mathbf{q}') \exp\left(-j2\pi\mathbf{q}'^T \mathbf{R}'_{k,l,m}\right) d\mathbf{q}' = \frac{1}{Q} P_\theta(\mathbf{R}'_{k,l,m}), \end{aligned} \quad (8)$$

where $Q = Q_x Q_y Q_z$ and $\mathbf{R}'_{k,l,m} = [k/Q_x, l/Q_y, m/Q_z]^T$ draws a regular lattice in the space of the EAP. Eqs. (7) and (8) can now be combined to establish the linear relation between the measurements in the q-space and the values of the EAP to estimate. Since \tilde{E}_θ is real and antipodal symmetric, a cosine series expansion may be used:

$$\tilde{E}_\theta(\mathbf{q}') = \frac{1}{Q} \sum_{k=-\infty}^{\infty} \sum_{l=-\infty}^{\infty} \sum_{m=-\infty}^{\infty} P_\theta(\mathbf{R}'_{k,l,m}) \cos\left(2\pi\mathbf{q}'^T \mathbf{R}'_{k,l,m}\right). \quad (9)$$

2.3. Building the encoding matrix

Eq. (9) will be used to estimate a Cartesian sampling of the EAP from a spherical multi-shell q-space sampling, $\{\mathbf{q}'_i \in \Omega\}_{i=1}^{N_i}$. Since the available number of q-samples is limited, so it is the number of lattice nodes we can actually estimate for the EAP, and in practice the Fourier series expansion will be cropped to its first few coefficients assuming \tilde{E}_θ is smooth enough:

$$\tilde{E}_\theta(\mathbf{q}') \simeq \frac{1}{Q} \sum_{k=-N_x}^{N_x} \sum_{l=-N_y}^{N_y} \sum_{m=-N_z}^{N_z} P_\theta(\mathbf{R}'_{k,l,m}) \cos\left(2\pi\mathbf{q}'^T \mathbf{R}'_{k,l,m}\right), \quad (10)$$

so that $2N_x + 1$ samples evenly spaced $1/Q_x$ should cover the entire 'x' axis of the EAP domain; $2N_y + 1$ samples, one each $1/Q_y$, should cover the 'y' axis; $2N_z + 1$ samples, one each $1/Q_z$, should cover the 'z' axis. This means that the compact support of E_θ is constrained by the number of lattice nodes along each dimension, which in turn is related to the available number of q-space samples, N_i . In precise terms:

- If the EAP is roughly compact supported at (i. e. it has negligible values outside of) the 3-D domain $[-\frac{X}{2}, \frac{X}{2}] \times [-\frac{Y}{2}, \frac{Y}{2}] \times [-\frac{Z}{2}, \frac{Z}{2}]$, sampling its entire support will demand:

$$\begin{aligned} \frac{N_x}{Q_x} &= \frac{X}{2} \Rightarrow Q_x = \frac{2N_x}{X}; & \frac{N_y}{Q_y} &= \frac{Y}{2} \Rightarrow Q_y = \frac{2N_y}{Y}; \\ \frac{N_z}{Q_z} &= \frac{Z}{2} \Rightarrow Q_z = \frac{2N_z}{Z}, \end{aligned} \quad (11)$$

so that any q-samples with either $|q_x| > \frac{Q_x}{2}$, $|q_y| > \frac{Q_y}{2}$, or $|q_z| > \frac{Q_z}{2}$ will be discarded since they must be assumed to be 0. Obviously, the larger N_x , N_y , and N_z , the fewer q-samples will need to be discarded. However:

- $N_i = (2N_x + 1)(2N_y + 1)(2N_z + 1)$ should be in the same order as N_i so that the number of parameters to estimate is not vastly greater than the number of measurements (in the absence of regularization, $N_i \leq N_i$ indeed).

Fortunately, the antipodal symmetry of the EAP ($P_\theta(\mathbf{R}') = P_\theta(-\mathbf{R}')$) allows to nearly halve the number of Degrees of Freedom (DoF) to estimate. From Eq. (10), a simple reordering of the addends yields:

$$\begin{aligned} E_\theta(\mathbf{q}') &\simeq \frac{1}{Q} \left(P_\theta(\mathbf{0}) + 2 \sum_{k=1}^{N_x} P_\theta(\mathbf{R}'_{k,0,0}) \cos\left(2\pi\mathbf{q}'^T \mathbf{R}'_{k,0,0}\right) \right. \\ &\quad \left. + 2 \sum_{k=-N_x}^{N_x} \sum_{l=1}^{N_y} P_\theta(\mathbf{R}'_{k,l,0}) \cos\left(2\pi\mathbf{q}'^T \mathbf{R}'_{k,l,0}\right) \right) \end{aligned}$$

$$+ 2 \sum_{k=-N_x}^{N_x} \sum_{l=-N_y}^{N_y} \sum_{m=1}^{N_z} P_\theta(\mathbf{R}'_{k,l,m}) \cos\left(2\pi\mathbf{q}'^T \mathbf{R}'_{k,l,m}\right). \quad (12)$$

Hence, the actual number of DoF to estimate is $n = 1 + N_x + (2N_x + 1)N_y + (2N_x + 1)(2N_y + 1)N_z = (N_i + 1)/2$. Now, let us substitute the triple indexing $\{k, l, m\}$ of the lattice with a unique index j by simply stacking its nodes in order:

$$\begin{aligned} \{\mathbf{R}'_j\}_{j=1}^n &= \{\mathbf{R}'_1, \mathbf{R}'_2, \dots, \mathbf{R}'_n\} \\ &\equiv \left\{ \mathbf{R}'_{0,0,0}, \dots, \mathbf{R}'_{N_x,0,0}, \mathbf{R}'_{0,1,0}, \dots, \mathbf{R}'_{N_x,1,0}, \dots, \mathbf{R}'_{N_x,N_y,0}, \right. \\ &\quad \left. \mathbf{R}'_{-N_x,-N_y,1}, \dots, \mathbf{R}'_{N_x,-N_y,1}, \dots, \mathbf{R}'_{N_x,N_y,1}, \dots, \mathbf{R}'_{N_x,N_y,N_z} \right\}. \end{aligned} \quad (13)$$

The $N_i \times n$ encoding matrix F relates the $N_i \times 1$ vector of measurements, $\mathbf{E} = [E_\theta(\mathbf{q}'_1), \dots, E_\theta(\mathbf{q}'_{N_i})]^T$, with the $n \times 1$ vector of unknowns, $\mathbf{P} = [P_\theta(\mathbf{R}'_1), \dots, P_\theta(\mathbf{R}'_n)]^T$. After Eq. (12):

$$\mathbf{E} \simeq F \mathbf{P} : [F]_{ij} = \frac{\kappa_j}{Q} \cos\left(2\pi\mathbf{q}'_i^T \mathbf{R}'_j\right), \quad (14)$$

where $\kappa_j = 1$ if $j = 0$ or $\kappa_j = 2$ otherwise. Finally, it is worth noticing the actual number of samples N_i may be voxel-dependent in case certain q-space samples lay outside the allowed bandwidth defined by Q_x , Q_y , and Q_z at each voxel.

2.4. Estimating the EAP from the encoding matrix

Eq. (14) establishes a linear relation between the vector of measurements \mathbf{E} and the vector of unknowns \mathbf{P} that allows solving for the latter with linear Least Squares (LS) techniques. However, two additional requirements must be fulfilled for the solution to be physically meaningful: positivity and unit mass. For the former, we will constrain all the entries of \mathbf{P} to be non-negative. For the latter, we note that the unit mass property of the EAP is equivalent to the attenuation signal evaluating to 1 at the origin $\mathbf{q}'_0 = \mathbf{0}$. Thus, the following QP problem arises:

$$\min_{\mathbf{P}} \frac{1}{2} \|\mathbf{E} - F\mathbf{P}\|^2 \text{ s. t. } \mathbf{P} \geq 0 \text{ and } f_0^T \mathbf{P} = 1, \quad (15)$$

where the $N_i \times 1$ vector $f_0 = \frac{1}{Q}[1, 2, \dots, 2]^T$ stands for an additional row of the encoding matrix at $\mathbf{q}'_0 = \mathbf{0}$. Nonetheless, these reconstruction problems often require some sort of regularization: first, the measurements vector \mathbf{E} is highly contaminated with noise. This is addressed by Wu et al. (2008), Wu and Alexander (2007) by just dropping down to 0 those values of the diffusion weighted images below a certain threshold (usually, twice the free-air average signal). In our case, we cannot forget that the actual number of samples, N_i , is voxel-dependent since out-of-bandwidth samples are removed, so that the QP in Eq. (15) might even become ill-posed. For these reasons, the QP problem will be reformulated as follows:

$$\min_{\mathbf{P}} \frac{1}{2} \|\mathbf{E} - F\mathbf{P}\|^2 + \frac{\lambda}{2} \|\mathcal{L}\mathbf{P}\|^2 \text{ s. t. } \mathbf{P} \geq 0 \text{ and } f_0^T \mathbf{P} = 1, \quad (16)$$

where $\lambda > 0$ is a small positive constant and \mathcal{L} is some linear operator. We will resort to a Laplacian penalty, described in the Appendix, as a popular choice to promote the smoothness of the solution (Descoteaux et al., 2007; Caruyer and Deriche, 2012; Fick et al., 2016b). Note Eq. (16) describes a convex problem, so that a unique optimum exists (Luenberger and Ye, 2008).

3. Computation of diffusion markers from the Cartesian EAP

Once the EAP is fully sampled in its whole domain, any numerical feature at will can be estimated from it. In this section we derive expressions for several commonly used diffusion markers: the RTOP and MSD, like Wu and Alexander (2007) do, but also the RTAP and the RTPP.

3.1. RTOP

Since the RTOP (or Po) is defined as the value of the EAP at the origin, it may be trivially computed as:

$$\text{RTOP} = P(\mathbf{0}) = P_{\Theta}(\mathbf{0}) = [\mathbf{P}]_1, \quad (17)$$

i.e. as the first component of vector \mathbf{P} .

3.2. RTAP

The RTAP represents the probability of water molecules moving back to the axis following the maximum diffusion direction within a time τ . In our model, such axis reduces to 'z'. The RTAP can be defined either on the EAP domain or the q-space:

$$\begin{aligned} \text{RTAP} &= \int_{-\infty}^{\infty} P_{\Theta}(\mathbf{R}') dz = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} E_{\Theta}(\mathbf{q}') dq_x dq_y \\ &= \int_{-Q_x/2}^{Q_x/2} \int_{-Q_y/2}^{Q_y/2} \tilde{E}_{\Theta}(\mathbf{q}') dq_x dq_y. \end{aligned} \quad (18)$$

The latter expression can be used to derive the RTAP from Eq. (12):

$$\begin{aligned} \text{RTAP} &\simeq \frac{1}{Q} \left(\int_{-\frac{Q_x}{2}}^{\frac{Q_x}{2}} \int_{-\frac{Q_y}{2}}^{\frac{Q_y}{2}} P_{\Theta}(\mathbf{0}) dq_x dq_y \right. \\ &+ 2 \sum_{k=1}^{N_x} P_{\Theta}(\mathbf{R}'_{k,0,0}) \int_{-\frac{Q_x}{2}}^{\frac{Q_x}{2}} \int_{-\frac{Q_y}{2}}^{\frac{Q_y}{2}} \cos(2\pi \mathbf{q}'^T \mathbf{R}'_{k,0,0}) dq_x dq_y \\ &+ 2 \sum_{k=-N_x}^{N_x} \sum_{l=1}^{N_y} P_{\Theta}(\mathbf{R}'_{k,l,0}) \int_{-\frac{Q_x}{2}}^{\frac{Q_x}{2}} \int_{-\frac{Q_y}{2}}^{\frac{Q_y}{2}} \cos(2\pi \mathbf{q}'^T \mathbf{R}'_{k,l,0}) dq_x dq_y \\ &+ 2 \sum_{k=-N_x}^{N_x} \sum_{l=-N_y}^{N_y} \sum_{m=1}^{N_z} P_{\Theta}(\mathbf{R}'_{k,l,m}) \int_{-\frac{Q_x}{2}}^{\frac{Q_x}{2}} \int_{-\frac{Q_y}{2}}^{\frac{Q_y}{2}} \cos(2\pi \mathbf{q}'^T \mathbf{R}'_{k,l,m}) dq_x dq_y \left. \right) \\ &= \frac{1}{Q_z} \left(P_{\Theta}(\mathbf{0}) + 2 \sum_{m=1}^{N_z} P_{\Theta}(\mathbf{R}'_{0,0,m}) \right). \end{aligned} \quad (19)$$

Remarkably, Eq. (19) equals the first order quadrature for the first integral form in Eq. (18), but it is exact up to the necessary cropping of the Fourier series coefficients.

3.3. RTPP

The RTPP represents the probability of water molecules moving back to the plane perpendicular to the maximum diffusion direction within a time τ . Since we assimilate the maximum diffusion direction to the 'z' axis, such domain is trivially described as the 'x'-'y' plane. The RTPP can be defined either on the EAP domain or the q-space:

$$\text{RTPP} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} P_{\Theta}(\mathbf{R}') dx dy = \int_{-\infty}^{\infty} E_{\Theta}(\mathbf{q}') dq_z = \int_{-Q_z/2}^{Q_z/2} \tilde{E}_{\Theta}(\mathbf{q}') dq_z. \quad (20)$$

The latter expression can be used to derive the RTPP from Eq. (12):

$$\begin{aligned} \text{RTPP} &\simeq \frac{1}{Q} \left(\int_{-\frac{Q_z}{2}}^{\frac{Q_z}{2}} P_{\Theta}(\mathbf{0}) dq_z + 2 \sum_{k=1}^{N_x} P_{\Theta}(\mathbf{R}'_{k,0,0}) \int_{-\frac{Q_z}{2}}^{\frac{Q_z}{2}} \cos(2\pi \mathbf{q}'^T \mathbf{R}'_{k,0,0}) dq_z \right. \\ &+ 2 \sum_{k=-N_x}^{N_x} \sum_{l=1}^{N_y} P_{\Theta}(\mathbf{R}'_{k,l,0}) \int_{-\frac{Q_z}{2}}^{\frac{Q_z}{2}} \cos(2\pi \mathbf{q}'^T \mathbf{R}'_{k,l,0}) dq_z \\ &+ 2 \sum_{k=-N_x}^{N_x} \sum_{l=-N_y}^{N_y} \sum_{m=1}^{N_z} P_{\Theta}(\mathbf{R}'_{k,l,m}) \int_{-\frac{Q_z}{2}}^{\frac{Q_z}{2}} \cos(2\pi \mathbf{q}'^T \mathbf{R}'_{k,l,m}) dq_z \left. \right) \\ &= \frac{1}{Q_x Q_y} \left(P_{\Theta}(\mathbf{0}) + 2 \sum_{k=1}^{N_x} P_{\Theta}(\mathbf{R}'_{k,0,0}) + 2 \sum_{k=-N_x}^{N_x} \sum_{l=1}^{N_y} P_{\Theta}(\mathbf{R}'_{k,l,0}) \right), \end{aligned} \quad (21)$$

which again equals the first order quadrature approximation to the first integral form in Eq. (20).

3.4. MSD

The MSD is the second order, non-central moment of the EAP. Since it will remain invariant to rotations, it can be equally computed from the rotated EAP $P_{\Theta}(\mathbf{R}')$:

$$\begin{aligned} \text{MSD} &= \iiint_{\mathbb{R}^3} \|\mathbf{R}'\|^2 P_{\Theta}(\mathbf{R}') d\mathbf{R}' \\ &= \iiint_{\mathbb{R}^3} (x^2 + y^2 + z^2) P_{\Theta}(\mathbf{R}') dx dy dz. \end{aligned} \quad (22)$$

From the theory on Fourier analysis, this quantity may be equally computed in the q-space by evaluating the (scaled) Laplacian of $E_{\Theta}(\mathbf{q}')$ at $\mathbf{q}' = \mathbf{0}$ (Oppenheim et al., 1997), as it has been thoroughly exploited in the literature (Tristán-Vega et al., 2009, 2010):

$$\text{MSD} = \frac{-1}{4\pi^2} \Delta_{\mathbf{q}'} E_{\Theta}(\mathbf{0}) = \frac{-1}{4\pi^2} \Delta_{\mathbf{q}'} \tilde{E}_{\Theta}(\mathbf{0}). \quad (23)$$

This expression allows a straightforward computation from Eq. (12):

$$\begin{aligned} \text{MSD} &\simeq \frac{-1}{4\pi^2 Q} \left(2 \sum_{k=1}^{N_x} P_{\Theta}(\mathbf{R}'_{k,0,0}) \Delta_{\mathbf{q}'} \cos(2\pi \mathbf{q}'^T \mathbf{R}'_{k,0,0}) \right. \\ &+ 2 \sum_{k=-N_x}^{N_x} \sum_{l=1}^{N_y} P_{\Theta}(\mathbf{R}'_{k,l,0}) \Delta_{\mathbf{q}'} \cos(2\pi \mathbf{q}'^T \mathbf{R}'_{k,l,0}) \\ &+ 2 \sum_{k=-N_x}^{N_x} \sum_{l=-N_y}^{N_y} \sum_{m=1}^{N_z} P_{\Theta}(\mathbf{R}'_{k,l,m}) \Delta_{\mathbf{q}'} \cos(2\pi \mathbf{q}'^T \mathbf{R}'_{k,l,m}) \left. \right) \Big|_{\mathbf{q}'=\mathbf{0}} \\ &= \frac{2}{Q} \left(\sum_{k=1}^{N_x} P_{\Theta}(\mathbf{R}'_{k,0,0}) \|\mathbf{R}'_{k,0,0}\|^2 + \sum_{k=-N_x}^{N_x} \sum_{l=1}^{N_y} P_{\Theta}(\mathbf{R}'_{k,l,0}) \|\mathbf{R}'_{k,l,0}\|^2 \right. \\ &+ \left. \sum_{k=-N_x}^{N_x} \sum_{l=-N_y}^{N_y} \sum_{m=1}^{N_z} P_{\Theta}(\mathbf{R}'_{k,l,m}) \|\mathbf{R}'_{k,l,m}\|^2 \right). \end{aligned} \quad (24)$$

Again, the first-order quadrature approximation to the integral in Eq. (22).

4. Numerical methods and algorithm parameters

4.1. Bandwidth selection

A key limitation of the original HYDI-DSI is the direct dependence of the bandwidth of the EAP with the particular q-space sampling scheme, which we address here. The lattice where the EAP will be defined will cover a domain that directly depends on the number of lattice points and its bandwidth (i.e. the support $\Omega \subset \mathbb{R}^3$ of the attenuation signal $E(\mathbf{q})$). For example, for the 'x' axis, the EAP will be sampled from $-N_x/Q_x$ to N_x/Q_x . The extent of the EAP domain actually covered can be increased by decreasing Q_x : this implies reducing the bandwidth of the signal, i.e. smoothing the EAP itself. Accordingly, the value of derived indices like RTOP will be artificially reduced due to the convolution of the EAP with a low-pass kernel. Conversely, we can keep a large bandwidth of the EAP by increasing Q_x , but in this case the maximum value sampled, N_x/Q_x , is likely not to cover a proper extent of the EAP. Fig. 2 illustrates these two issues. Of course, we can think of increasing N_x together with Q_x to get full EAP coverage while preserving and adequate bandwidth. But the number of lattice nodes, in the order of $\mathcal{O}(N_x \cdot N_y \cdot N_z)$, cannot be arbitrarily large, but instead it should roughly match the number of available q-space samples even if regularization constraints are imposed.

Since the lattice axes are aligned with the eigenvectors of the low b-value regimen, DTI approximation, it makes sense to scale them according to the respective eigenvalues. For the 'x' axis, the farthest sampled value can be chosen such that the Gaussian tail has decayed




Fig. 2. The effect of wrongly choosing the bandwidth of the EAP for $N_x = 4$: if Q_x is too small (red), the EAP support is fully covered but it becomes low-pass filtered; if Q_x is too large (green), the EAP is kept sharp but its support is not properly sampled.

to a pre-defined small value $\mu > 0$:

$$\begin{aligned}
 P_{\Theta}(0, 0, 0) &= \frac{\exp\left(-\frac{1}{4\tau\lambda_1}\left(\frac{0}{Q_x}\right)^2\right)}{\sqrt{\det(\Lambda)(4\pi\tau)^3}} \\
 P_{\Theta}\left(\frac{N_x}{Q_x}, 0, 0\right) &= \frac{\exp\left(-\frac{1}{4\tau\lambda_1}\left(\frac{N_x}{Q_x}\right)^2\right)}{\sqrt{\det(\Lambda)(4\pi\tau)^3}} = \mu P_{\Theta}(0, 0, 0) \\
 \Rightarrow \frac{Q_x}{2} &= \frac{N_x}{4\sqrt{-\tau\lambda_1 \log(\mu)}}.
 \end{aligned} \quad (25)$$

Note this expression can be otherwise written in terms of maximum b-values to get rid of the diffusion time τ , so that:

$$\begin{aligned}
 b_{\text{cut-off},x} &= 4\pi^2\tau\left(\frac{Q_x}{2}\right)^2 = \frac{-\pi^2 N_x^2}{4\lambda_1 \log(\mu)}; \\
 b_{\text{cut-off},y} &= \frac{-\pi^2 N_y^2}{4\lambda_2 \log(\mu)}; \quad b_{\text{cut-off},z} = \frac{-\pi^2 N_z^2}{4\lambda_3 \log(\mu)}.
 \end{aligned} \quad (26)$$

It becomes inherent to our approach that q-space samples measured beyond these limits have to be dropped down to avoid aliasing. Consequently, the parameter μ must be appropriately tuned. As a final remark, there is no reason why this same bandwidth tuning cannot be used with the original re-gridding/interpolation approach to define the domain Ω : it suffices to input zeros to the interpolation algorithm at boundary points $\{-Q_x/2, 0, Q_x/2\} \times \{-Q_y/2, 0, Q_y/2\} \times \{-Q_z/2, 0, Q_z/2\} - \mathbf{0}$, and then operate as described by Wu et al. (2008), Wu and Alexander (2007) for an input lattice defined over $\Omega = [-Q_x/2, Q_x/2] \times [-Q_y/2, Q_y/2] \times [-Q_z/2, Q_z/2]$ computed after Eq. (25). We will put this strategy to the test as described in Section 5.2.

4.2. Practical implementation

The computation of the Gaussian approximation in Eq. (5) is accomplished with linearized LS techniques from the q-space samples with $b \leq 2000$ s/mm² (Salvador et al., 2005). The bandwidth of the signal, described by $b_{\text{cut-off},e_j}$, will be determined according to Eq. (26) for an empirically fixed μ . The number of samples along each dimension is a design parameter, but we will use the same sampling as in the original HYDI-DSI approach, i. e. a $9 \times 9 \times 9$ lattice ($N_x = N_y = N_z = 4$) with $N_l = 729$, so that $n = (N_l + 1)/2 = 365$ DoF have to be estimated.

The form of matrix \mathcal{L} in Eq. (16), describing the Laplacian energy penalty, is described in the Appendix, and the parameter λ will be fixed according to empirical considerations.

Finally, the QP in Eq. (16) is solved with an *ad hoc* replacement of Matlab's quadprog function, based on gradient projection (Luenberger and Ye, 2008). The iterations are initialized with the unconstrained solution of Eq. (16), which reduces to the computation of $(F^T F + \lambda L^T L)^{-1} F^T E$, where the matrix to invert is symmetric and positive definite. Such solution is corrected for negative values and normalized to fulfill the unit-mass constraint before feeding the iterations

until convergence. Note the QP is convex, so it is always guaranteed to converge to the global optimum. Matlab code can be downloaded as a part of the *dMRI-Lab* toolbox (<http://www.lpi.tel.uva.es/dmriLab>).

5. Experiments and results

5.1. Materials

The evaluation of the proposal will be based on publicly available databases. Concretely:

- The **Human Connectome Project (HCP)**, MGH database (Fan et al., 2016): these are high quality data acquired on a Siemens 3T Connectome scanner with 4 different shells at $b = \{1000, 3000, 5000, 10,000\}$ s/mm², with $\{64, 64, 128, 256\}$ gradient directions each and 40 interleaved unweighted baselines. The in-plane resolution is 1.5 mm and the slice thickness is 1.5 mm. Other acquisition parameters are: pulse separation time/diffusion gradients length $\Delta/\delta = 21.8/12.9$ ms, repetition time TR = 8800 ms, time echo TE = 57 ms. We have randomly chosen subject HCP MGH-1007 for proofs of concept.
- The **HCP WU-Minn database** (Van Essen et al., 2013): these data were acquired with a Siemens 3T Connectome Skyra scanner with a maximum gradient strength at 100 mT/m, 3 shells at $b = \{1000, 2000, 3000\}$ s/mm² with 90 gradient directions each and 18 interleaved unweighted baselines. The in-plane resolution is 1.25 mm and the slice thickness is 1.25 mm. Other acquisition parameters are: $\Delta/\delta = 43/10.6$ ms, TR/TE = 5520/89.5 ms. We selected 10 subjects, enrolled for both test and re-test acquisition sessions: 103818, 105923, 111312, 114823, 115320, 122317, 125525, 130518, 139839, 143325.

5.2. Methods compared

The original method as described by Wu et al. (2008), Wu and Alexander (2007) will be simply referred to as HYDI-DSI. We will use a custom implementation based on Matlab's `griddatan` as the authors suggest. Note, however, that the q-space grid will be rotated at each voxel (but not stretched) before the interpolation, which allows us to align the maximum diffusion direction with the 'z' axis to compute RTAP and RTPP. The proposed method described in this paper will be dubbed HYDI-DSI-QP attending to its numerical implementation as a QP. Besides, we will probe a third method, HYDI-DSI-AB, for which q-space re-gridding and interpolation is used in a lattice defined over an adaptive bandwidth (AB) as described in Section 4.1. This way, HYDI-DSI-AB is an intermediate approach that inherits from both HYDI-DSI and HYDI-DSI-QP, and it can indeed be considered a novel method proposed in the present paper. Finally, we will include comparisons with MAPL (Fick et al., 2016b), which is probably the most popular approach within the category of continuous domain, parametric representations of the EAP. We will use Python's `dipy` implementation (<https://dipy.org/>) with anisotropic scaling, positivity constraints, a maximum radial order 4 for the basis functions, and a fixed value 0.2 for the Laplacian weighting. While order 6 is advised by the authors in case isotropic scaling is used, order 4 provides an acceptable trade-off between representation capabilities and computational load in the anisotropic case.

5.3. Bandwidth selection

The bandwidth chosen for the signal at each voxel depends on the parameter μ and the lattice size, but not on the particular q-space sampling scheme. Indeed, only those b-values below 2000 s/mm² are actually used to fit the DTI approximation. Accordingly, we have chosen subject HCP MGH-1007 as a representative example to elaborate Fig. 3, where the cut-off b-values at each (rotated) axis 'x', 'y', and

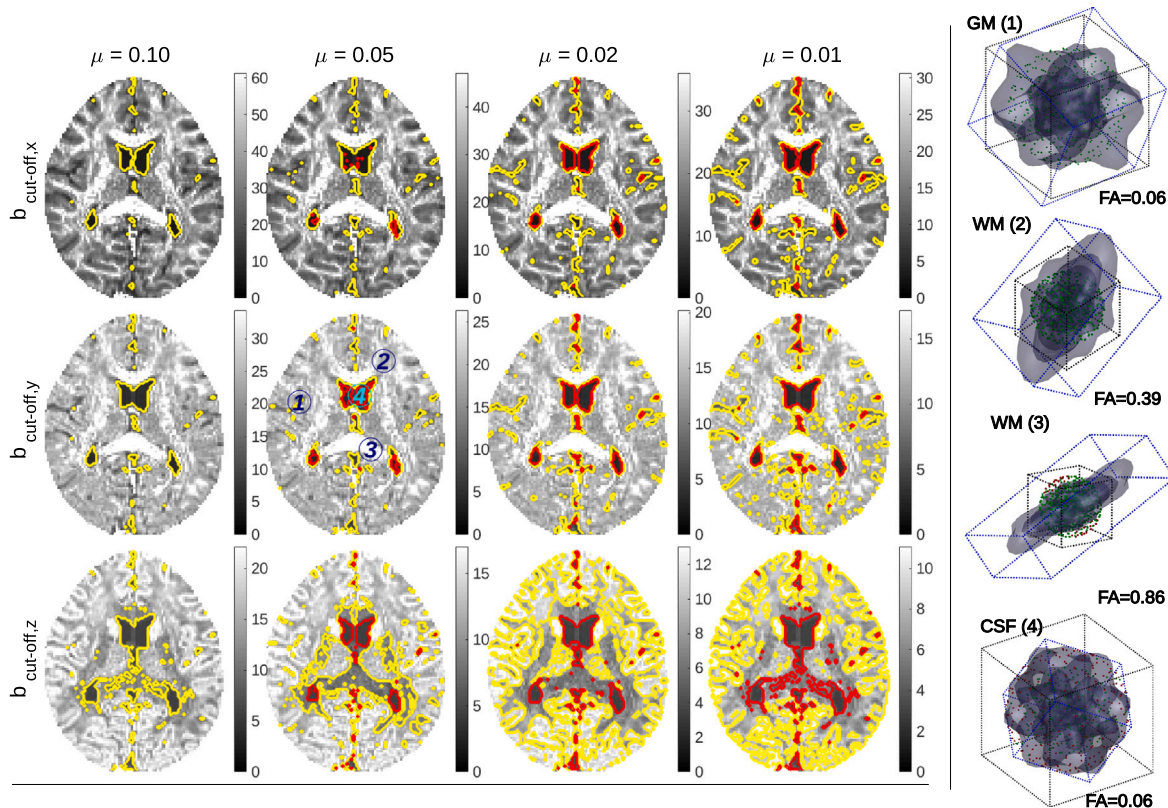


Fig. 3. (Left) Cut-off b-values ($\times 10^3$ s/mm²) for the three main diffusion directions (top to bottom, on ascending order of the eigenvalues) and for several values of parameter μ , all of them computed over volume HCP MGH-1007 for a fixed lattice size $9 \times 9 \times 9$. Yellow lines correspond to iso-contours at $b_{\text{cut-off},e_j} = 10,000$ s/mm², and red lines to iso-contours at $b_{\text{cut-off},e_j} = 5000$ s/mm². (Right) A detail of the bandwidth selection ($\mu = 0.05$) at four representative voxels located as shown in the left figure. The black bounding boxes represent the original q-space domain given by the maximum b-value 10,000 s/mm²; blue ones stand for the estimated bandwidths in the rotated space from the DTI approach. The dots represent the acquired b-values and gradients (red: discarded; green: used), and the surfaces depict the actual attenuation signal at each shell.

‘z’ are represented at a middle-brain axial slice. We have chosen a typical $9 \times 9 \times 9$ lattice with $N_x = N_y = N_z = 4$, corresponding to $n = 365$ DoF of the EAP. Fig. 3 (left) allows to conclude that, unless a extreme value is chosen for μ , such as 0.01, the only samples rejected at $b = 5000$ s/mm² and above correspond to the cerebrospinal fluid (CSF). To corroborate this assertion, Fig. 3 (right) shows a typical voxel at the CSF (4), for which the estimated bandwidth is smaller than the sampled bandwidth and a large number of samples are dropped out.

On the contrary, even for $\mu = 0.05$, many non-CSF samples at $b = 10,000$ s/mm² will be discarded for axis ‘z’. As it can be inferred from Fig. 3 (right, voxel 3), this situation corresponds to prominent, well packaged structures (with large Fractional Anisotropy, FA) within the white matter (WM), like the corpus callosum (CC; voxel 3) and the cortico-spinal tract (CST), for which the signal along the main diffusion direction rapidly vanishes: even when the estimated bandwidth is larger than the sampled one at the transverse diffusion plane, the bandwidth for the main diffusion direction crops the sampled one, so that the samples at the North and South poles of the rotated space are discarded. Note this artifact is likely to appear also in the original HYDI-DSI, since the signal at these points will fall below the threshold set as twice the free air signal.

With the same value $\mu = 0.05$, Fig. 3 (right, voxels 1 and 2), shows that no samples will be discarded, in general, at the gray matter (GM) or low-FA regions of the WM. For the latter, a low FA value will reflect regions with important partial volume effects due to fiber crossings and/or bending, for which any spatial direction will mix up both restricted and free diffusion.

Finally, $\mu = 0.10$ practically avoids out-of-bandwidth q-space pruning, but this comes at the expense of a poor coverage of the tails of the EAP. In the light of this experiment, we can conclude that applying a

threshold $\mu = 0.05$ to the tails of the EAP seems a good trade-off for most of situations.

5.4. Regularization parameter selection

The optimal value of the regularization parameter λ in Eq. (16) is likely to depend on several factors, such as the q-space sampling and the SNR of the data set, the lattice size, or the bandwidth chosen through parameter μ . For this reason, we have explored the same previous values of μ and two data sets with very different characteristics, concretely HCP MGH-1007 and HCP WuMinn-139839, to trace the L-curves in the respective Fig. 4(a) and (b). Besides, in order to keep a reasonable complexity, we have fixed the lattice size once again to the standard $9 \times 9 \times 9$. The L-curves plot the trade-off between the fitting residual $\|\mathbf{E} - \mathcal{F}\mathbf{P}\|^2$ and the Laplacian energy penalty $\|\mathcal{L}\mathbf{P}\|^2$ as parameter λ varies: for large values of λ , the Laplacian penalty becomes stronger and smoother solutions are promoted at the expense of larger residuals. Conversely, for small values of λ the solution will resemble closer the acquired data, but it will likely become physically little plausible. We have focused on the white matter, which has been roughly segmented by thresholding the FA of the DTI approximation at 0.3. First thing to note is that larger bandwidths (i.e. larger μ) translate in smaller residuals, even when less q-space samples are discarded in the QP of Eq. (16) and hence vector \mathbf{E} has more components. This is an additional reason to avoid an excessive cropping of the large b-value regimen. Next, it seems the behavior of the L-curves is very little dependent on the value of μ , which allows to tune both parameters independently. Finally, the optimal value of λ can be assured to be in the range [0.1, 2.5] in all cases: the corner of the L-curves marks the point where a slight improvement in the residual will come at the expense of an important

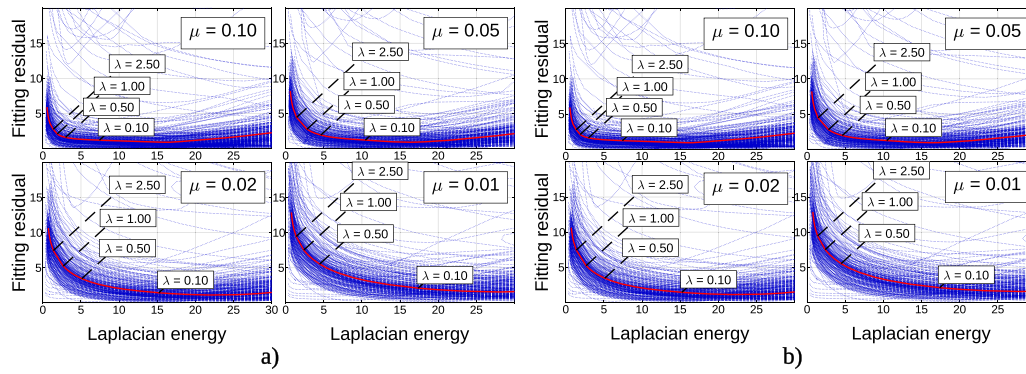


Fig. 4. Fitting residual $\|E - FP\|^2$ vs. Laplacian energy penalty $\|LP\|^2$ as a function of the regularization parameter λ , for subjects HCP MGH-1007 (a) and HCP WuMinn-139839 (b) and for several values of μ . The curves represented in dashed, blue lines correspond to a random sub-sample inside the white matter. The red curve in solid line represents the average value inside the white matter.

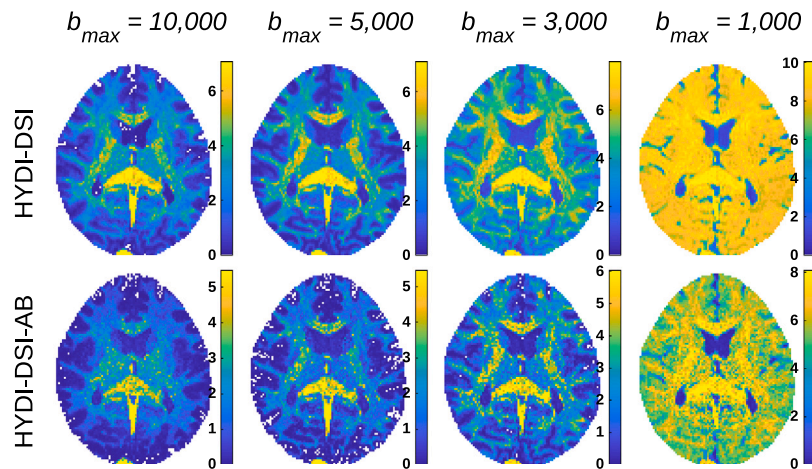


Fig. 5. Percentage of energy corresponding to negative values of the EAP within each voxel for the unconstrained methods: the original HYDI-DSI (top row) and the adaptive bandwidth method, HYDI-DSI-AB (bottom row).

worsening of the smoothness of the solution (and vice-versa), so the optimal λ should be chosen close to this point. Accordingly, a fixed value $\lambda = 0.50$ seems adequate for both data sets HCP MGH and HCP WuMinn, and it is the fixed value used throughout the paper.

5.5. Negativity of unconstrained methods

One main advantage of the HYDI-DSI-QP is the possibility it grants to enforce sampled EAP values to be positive, which is unparalleled for the interpolated methods (HYDI-DSI and HYDI-DSI-AB). Fig. 5 is aimed at checking the actual impact of such constraints depending on the number of acquired shells. Once again, subject HCP MGH-1007 has been considered as a representative example: both HYDI-DSI and HYDI-DSI-AB (with $\mu = 0.05$, according to the previous experiment) have been used to compute the EAP at each voxel inside $9 \times 9 \times 9$ lattices, and the percentage of energy corresponding to negative values of the EAP estimated by numerical quadrature. As it can be seen, the presence of negative lobes within the EAP is not a negligible issue. Indeed, it becomes more noticeable as the number of acquired shells decreases (as expected), but also it is more relevant for those white matter regions with highest anisotropy: in particular, nearly 10% (or more) of the estimated values of the EAP can be negative inside the CC and the CST for any number of acquired shells. Comparing HYDI-DSI with HYDI-DSI-AB, the adaptive bandwidth selection for the latter helps palliating this artifact to some degree. Finally, note the computed percentage of *negative energy* with HYDI-DSI-QP will always be zero by construction.

5.6. Accuracy of scalar maps

The quantitative evaluation of dMRI techniques is usually a challenging problem due to the lack of realistic ground-truth data. In this paper we will reuse the original approach conceived by Tristán-Vega and Aja-Fernández (2021), where an actual micro-structural model is estimated using the NODDI technique by Zhang et al. (2012) at representative regions of the white matter. The parameters obtained are statistically characterized to draw random samples that are fed to the forward NODDI model to generate synthetic samples simulating 1, 2, or 3 crossing fibers at will with known Peak Signal to Noise Ratio (PSNR). As long as the generative model can be sampled for any arbitrary gradient direction and b-value, ground-truth values are easily obtained for any dMRI measure with arbitrary precision by numerical integration. The reader is addressed to Tristán-Vega and Aja-Fernández (2021) for further details on this methodology.

Fig. 6 shows Bland & Altman agreement plots (Bland and Altman, 1986) between the indices (one of RTOP, RTAP, RTPP or MSD) as computed with either of the methods compared (one of MAPL, HYDI-DSI, HYDI-DSI-AB, or HYDI-DSI-QP) and the ground truth. Three different PSNR values were probed, and the plots comprise all possible fiber-crossing configurations in all cases. The q-space sampling scheme is the one found in the HCP MGH database. Since each method will potentially introduce a variable scale-shift for these indices depending on the estimated bandwidth, we use logarithmic plots, i.e.: if the agreement between M_1 and M_2 has to be assessed, we plot y vs. x , where x and y are respectively defined as $x = (\log(M_1) + \log(M_2))/2$




Fig. 6. Bland & Altman agreement plots (logarithmic) between the scalar measures obtained with each of the methods compared (M_1) and the synthetic ground truth (M_2), as a function of the PSNR and for mixed fiber configurations. The scalar measures tested are (a) RTOP, (b) RTAP, (c) RTPP and (d) MSD. The ‘x’ axis represents $(\log(M_1) + \log(M_2))/2$, and the ‘y’ axis $\log(M_2) - \log(M_1)$. Dashed lines represent median values, whereas dotted lines represent the 10% and 90% percentiles; the correlation coefficient ρ and the Normalized Mutual Information (NMI) are shown for quantitative assessment.

and $y = \log(M_2) - \log(M_1)$. For quantitative assessment, we provide also estimates of the correlation coefficient and the Normalized Mutual Information (NMI) between x_1 and x_2 in all cases.

In general terms, all methods attain fairly good results for medium-large PSNR, but their performance obviously worsens as the PSNR decreases. Comparing the proposed HYDI-DSI-QP with the original HYDI-DSI, the former is outperformed, both in terms of correlation and NMI, by the latter for the RTOP and RTAP, but it does overall better with the RTPP and MSD. HYDI-DSI-AB, in turn, does not seem consistently beneficial for any index. Yet, the original HYDI-DSI proves itself extraordinarily robust to noise for RTOP and RTAP, with the plots experiencing very little spreading as the PSNR decreases.

If we compare now with MAPL, its behavior is in general similar to that of HYDI-DSI-QP, though the former performs slightly better than the latter in some scenarios (above all, for RTAP and RTPP) in terms of correlations and NMI. Note, however, that the plots for HYDI-DSI-QP

are closer to 0 for all indices and all PSNR values, meaning the scale-shifts induced by our proposed technique are less important than those with all the other methods (except, in some situations, for HYDI-DSI). With regard to this issue, note that MAPL is indeed the worst performer, since it yields to a larger bias than all other methods in all situations.

5.7. Sensitivity of scalar maps to the number of acquired shells

One of the aims of the present proposal is reducing the dependency of quantitative dMRI parameters on the q-space sampling scheme by designing a sampling-independent bandwidth. To check this property, we compute the RTOP, RTAP, RTPP and MSD for subject HCP MGH-1007. At first instance, we will use the four available shells up to $b = 10,000 \text{ s/mm}^2$, since the original HYDI-DSI was designed for this range of maximum b-values. This estimate will be set as the bronze standard for each method, and then the outermost shells will be sub-




Fig. 7. RTOP values (in mm^{-3}) for subject HCP MGH-1007 using either the four available shells ($b_{\text{max}} = 10,000 \text{ s/mm}^2$), the three innermost shells ($b_{\text{max}} = 5000 \text{ s/mm}^2$), the two innermost shells ($b_{\text{max}} = 3000 \text{ s/mm}^2$), or just the innermost shell ($b_{\text{max}} = 1000 \text{ s/mm}^2$). Left: an illustrative central axial slice. Right: Bland & Altman agreement plots (logarithmic) between the computation with all four shells (x-axis) and the computation with the innermost shells (y-axis) as indicated. The correlation coefficient ρ and the NMI are also shown for quantitative assessment.




Fig. 8. RTAP values (in mm^{-2}) for subject HCP MGH-1007 using either the four available shells ($b_{\text{max}} = 10,000 \text{ s/mm}^2$), the three innermost shells ($b_{\text{max}} = 5000 \text{ s/mm}^2$), the two innermost shells ($b_{\text{max}} = 3000 \text{ s/mm}^2$), or just the innermost shell ($b_{\text{max}} = 1000 \text{ s/mm}^2$). Left: an illustrative central axial slice. Right: Bland & Altman agreement plots (logarithmic) between the computation with all four shells (x-axis) and the computation with the innermost shells (y-axis) as indicated. The correlation coefficient ρ and the NMI are also shown for quantitative assessment.

sequentially removed to estimate the same scalar measurements from either three, two, or even one shell.

Fig. 7 shows the results for RTOP, both qualitatively (RTOP maps, left) and quantitatively (Bland & Altman plots, right). As it can be seen, the benefit of HYDI-DSI-QP over all other methods compared is now clear: the plots are very well clustered around the median value in all cases demonstrating an outstanding agreement even if only one shell at $b_{\text{max}} = 1000 \text{ s/mm}^2$ is used, whereas all other methods (especially MAPL) yield widespread cloud points. This conclusion is supported by the fairly larger NMI value attained by HYDI-DSI-QP in all cases (the correlation coefficient is not a conclusive performance index since it becomes very close to 1 in all cases).

With regard to the scale shifts, HYDI-DSI is extremely sensitive to the removal of the outermost shells (because the estimated bandwidth

directly depends on the maximum b -value itself), meanwhile HYDI-DSI-AB and HYDI-DSI-QP are more robust. Though the estimated bandwidth for these two methods is acquisition protocol-independent, they will be affected by the lack of samples at large b -values as well: with HYDI-DSI-QP, the Laplacian penalty promotes smooth solutions; in the absence of high frequency q-samples to fit, the obvious way in which this condition is fulfilled is producing low-pass responses, which translates in reduced values of RTOP, see Fig. 2; with HYDI-DSI-AB, the missing samples are just replaced with a zero padding, which leads also to an artificial bandwidth drift. In this sense, the scales of the parameters estimated with MAPL remain more faithful to those at $b_{\text{max}} = 10,000 \text{ s/mm}^2$ (though, after Fig. 6, these scales might be distorted).

Very similar comments arise for the RTAP from Fig. 8: HYDI-DSI-QP is able to faithfully reproduce the original results with all four




Fig. 9. RTPP values (in mm^{-1}) for subject HCP MGH-1007 using either the four available shells ($b_{\max} = 10,000 \text{ s/mm}^2$), the three innermost shells ($b_{\max} = 5,000 \text{ s/mm}^2$), the two innermost shells ($b_{\max} = 3,000 \text{ s/mm}^2$), or just the innermost shell ($b_{\max} = 1,000 \text{ s/mm}^2$). Left: an illustrative central axial slice. Right: Bland & Altman agreement plots (logarithmic) between the computation with all four shells (x-axis) and the computation with the innermost shells (y-axis) as indicated. The correlation coefficient ρ and the NMI are also shown for quantitative assessment.

shells even if the sampling scheme is reduced to just one shell, which translates in well clustered plots and noticeably larger NMI values. Yet, HYDI-DSI-QP is able to better preserve the scale of the RTAP across all the experiments as compared to HYDI-DSI and even MAPL, being slightly outperformed only by HYDI-DSI-AB.

For the RTPP, Fig. 9 demonstrates that HYDI-DSI-QP is still the best performer, though its actual advantage over the other methods is not equally clear in this case. Besides, and conversely to the RTAP, the scale-shifts induced in the RTPP as the outermost shells are removed is especially noticeable with HYDI-DSI-QP, and only the original HYDI-DSI remains more sensitive. It is worth noticing that the Bland & Altman map for HYDI-DSI-AB at $b_{\max} = 5,000 \text{ s/mm}^2$ presents an artifact in the form of a spread cloud over-imposed to a clearly defined linear cluster. The RTPP can be computed as the integral of $E(\mathbf{q})$ along the main diffusion direction: for those voxels with a large FA, corresponding to a unique, well packaged, fiber bundle (e.g. voxel 3 in Fig. 3, right), the signal will very quickly decay for this main direction, so that it will very likely have vanished at $b = 5,000 \text{ s/mm}^2$ and above. Since HYDI-DSI-AB works by zero-padding all the values above b_{\max} , the estimates for $b_{\max} = 5,000 \text{ s/mm}^2$ and $b_{\max} = 10,000 \text{ s/mm}^2$ will be exactly the same, which explains the straight line in the figure. For other voxels with smaller FA, the estimates will largely differ due to the zero-padding, which explains the spread cloud.

Finally, the MSD is tested in Fig. 10 with similar conclusions: HYDI-DSI-QP provides very well clustered results, demonstrating an outstanding agreement with the whole sampling even if just one shell is used. The NMI values, once again, corroborate the visual inspection. With regard to the scale shifts, in this case HYDI-DSI-AB seems particularly robust, whereas HYDI-DSI-QP and MAPL behave similarly and HYDI-DSI, again, proves itself especially sensitive.

5.8. Repeatability and reliability analyses

The importance of computing quantitative indices in a reliable manner within neurosciences has been stressed by Zuo et al. (2019). According to the authors, such indices not only need to be repeatable for the same subject/anatomy being imaged (i.e. be specific), but they should also reflect meaningful anatomical changes among subjects (i.e. be sensitive). The quantitative assessment of repeatability through test–retest acquisitions has become a common topic in dMRI,

to the point that several *ad hoc* databases have been designed with this aim (Koller et al., 2021; Van Essen et al., 2013). In particular, we will use the test–retest subset of the HCP WuMinn database in this section.

We calculate the RTOP, RTAP, RTPP and MSD with MAPL, HYDI-DSI, HYDI-DSI-AB and HYDI-DSI-QP. Each combination involves both the test and the retest acquisitions within the data set. All three available shells and all 90 gradient directions per each shell were used. Besides, we estimate the diffusion tensors at $b = 1,000 \text{ s/mm}^2$ using the FSL 6.0.4 `dtifit` tool (Smith et al., 2004, Analysis Group, FMRIB, Oxford, UK; <https://fsl.fmrib.ox.ac.uk/fsl/fslwiki>), retrieve the FA parameter for all test/retest cases, and non-linearly register the FA maps using a normalized correlation cost function to the common template FMRIB58 (a high-resolution FA volume aggregated from 58 subjects) with a voxel resolution of $1 \times 1 \times 1 \text{ mm}^3$ (Jenkinson et al., 2002; Jenkinson and Smith, 2001). Eventually, all the above-mentioned propagator-based measures are non-linearly warped to the common space with trilinear interpolation. Once the measures are warped to the standard space, we calculate their reliability as:

$$\text{Reliability}(\mathbf{x}) = \frac{\text{Separability}(\mathbf{x})}{\text{Repeatability}(\mathbf{x})}, \quad (27)$$

with \mathbf{x} -dependent $\text{Repeatability}(\mathbf{x})$ averaged from the repeatabilities of each of the m out of M subjects:

$$\text{Repeatability}_m(\mathbf{x}) = \frac{\text{std. dev. across sessions}_s(\mathbf{x})}{\text{mean across sessions}_s(\mathbf{x})} \text{ for } s = \{\text{test, retest}\}. \quad (28)$$

$\text{Separability}(\mathbf{x})$ is conversely given by:

$$\text{Separability}(\mathbf{x}) = \frac{\text{std. dev. across subjects}_m(\text{mean across sessions}_s(\mathbf{x}))}{\text{mean across subjects}_m(\text{mean across sessions}_s(\mathbf{x}))} \text{ for } m = \{1, \dots, M\}. \quad (29)$$

Table 1 presents the results of the repeatability and reliability studies using ten subjects ($M = 10$) from the HCP WuMinn test–retest subset. In both cases, the numbers were averaged from 20 representative slices in the standard space and the WM area extracted from the John Hopkins University DTI-based atlas (Mori et al., 2005). Table 1(a) displays the coefficients of variation of the measures expressed as a

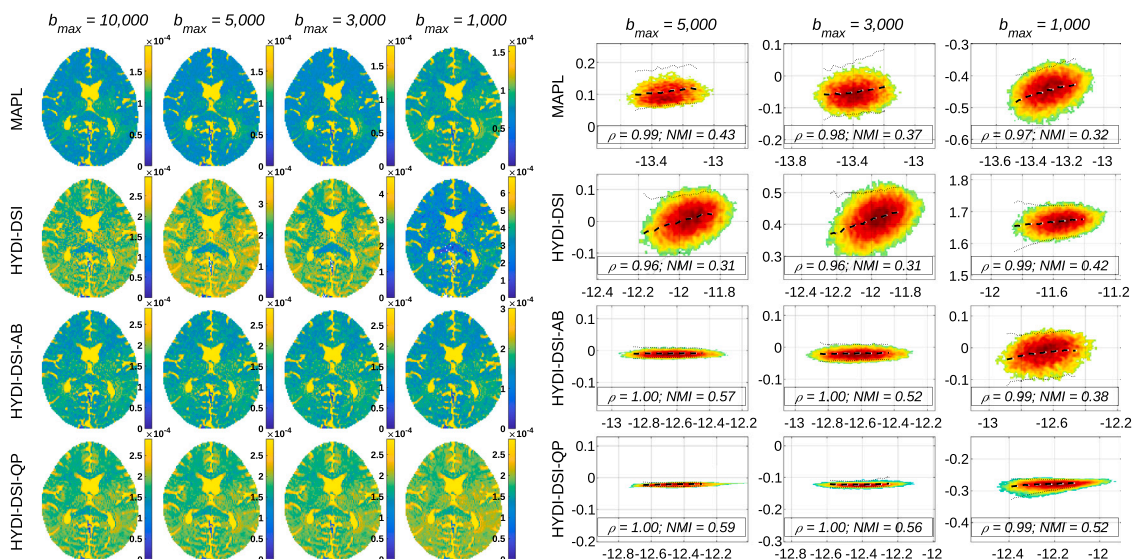


Fig. 10. MSD values (in mm^2) for subject HCP MGH-1007 using either the four available shells ($b_{\text{max}} = 10,000 \text{ s/mm}^2$), the three innermost shells ($b_{\text{max}} = 5000 \text{ s/mm}^2$), the two innermost shells ($b_{\text{max}} = 3000 \text{ s/mm}^2$), or just the innermost shell ($b_{\text{max}} = 1000 \text{ s/mm}^2$). Left: an illustrative central axial slice. Right: Bland & Altman agreement plots (logarithmic) between the computation with all four shells (x-axis) and the computation with the innermost shells (y-axis) as indicated. The correlation coefficient ρ and the NMI are also shown for quantitative assessment.

Table 1

Averaged repeatability (a) and reliability (b) inside the white matter from the test–retest subset of the HCP WuMinn database. The smaller, the better reproducibility, while the higher, the better reliability.

(a)				
Repeatability (in %)	RTOP	RTAP	RTPP	MSD
MAPL	6.31	5.15	1.90	4.38
HYDI-DSI	2.12	1.98	1.67	4.24
HYDI-DSI-AB	8.18	6.85	2.36	3.52
HYDI-DSI-QP	4.78	3.76	1.45	3.04

(b)				
Reliability	RTOP	RTAP	RTPP	MSD
MAPL	2.79	3.12	3.10	2.45
HYDI-DSI	2.73	3.07	2.23	2.49
HYDI-DSI-AB	2.79	3.05	2.47	2.38
HYDI-DSI-QP	2.73	3.05	3.47	2.84

percentage (i.e. multiplied by 100). This means the smaller the repeatability score, the better the reproducibility achieved. Generally, all measures are characterized by a high level of reproducibility, with the RTPP being the most reproducible measure (the coefficient of variation is at most 2.36%). The HYDI-DSI-QP technique achieves improved results over HYDI-DSI-AB including all measures and is superior to the standard HYDI-DSI technique for RTPP and MSD. However, in the case of RTOP and RTAP, the HYDI-DSI provides better performance over both HYDI-DSI-AB and HYDI-DSI-QP. Noticeably, both HYDI-DSI and HYDI-DSI-QP consistently outperform MAPL for all indices with regard to repeatability. The averaged reliability of the measures, as defined in Eq. (27), is presented in Table 1(b). In this experiment, the higher the value, the better the reliability. The results show comparable (virtually identical) reliability for RTOP and RTAP among all four methods, despite large discrepancies occur if we pay attention only to the repeatability. For RTPP and MSD, HYDI-DSI-QP is notably superior to all the other methods compared.

5.9. Execution times

Quantitative dMRI is often computationally very demanding, entailing processing times that range from several minutes to many hours or

Table 2

Average per-voxel execution times (milliseconds) depending on the number of shells considered.

	4 shells	3 shells	2 shells	1 shell
MAPL	368.2	303.1	252.1	203.7
HYDI-DSI	17.6	4.8	10.3	2.8
HYDI-DSI-AB	14.3	9.4	4.9	4.7
HYDI-DSI-QP	3.0	1.8	1.1	0.9

even days per subject. Hence, the study of the computational complexity is undoubtedly interesting in this context. Table 2 summarizes the average computation times taken by each method compared in this paper to process one single voxel. We have used subject HCP MGH-1007 for illustration purposes, and tested several sampling schemes with the same fixed configurations described in Section 5.7: four shells up to $b = 10,000 \text{ s/mm}^2$, three shells up to $b = 5000 \text{ s/mm}^2$, two shells up to $b = 3000 \text{ s/mm}^2$, or one shell at $b = 1000 \text{ s/mm}^2$. Times are reported for Matlab R2020a implementations (HYDI-DSI-like) or Python 3.6 (MAPL) running on an Intel(R) Xeon(R) E5-2695 CPU with 54 cores at 2.30 GHz and 110 GB RAM memory under Ubuntu Linux 20.04, all of them multi-threaded. As it could be expected beforehand, MAPL and HYDI-DSI-QP are always faster as the number of q-space samples decreases, since the size of the QP is obviously smaller. This comment does not hold for HYDI-DSI, for which the bottleneck is in the computation of the convex hull and Delaunay’s triangulation (linear interpolation will take a negligible time): the computation time is not always monotonically increasing with the number of samples, which makes this method less predictable with regard to its complexity. Though HYDI-DSI-AB takes decreasing times with smaller samplings, it is still prone to the same issue as HYDI-DSI. HYDI-DSI and HYDI-DSI-AB are comparable in all cases, each one being faster for certain configurations and slower for others. In any case, the proposed HYDI-DSI-QP is notably faster than the re-gridding/interpolation methods, with an acceleration factor ranging $3\times$ to $10\times$. Finally, MAPL is dramatically slower than all other methods, nearly $100\times$ to $200\times$ compared to HYDI-DSI-QP: meanwhile a complete volume from the HCP database can be processed with the latter in roughly half an hour, the former will take well over two days.

6. Discussion and conclusions

The proposed HYDI-DSI-QP is able to estimate fully non-parametric, positive, unit-mass constrained EAPs at a regular Cartesian lattice from arbitrary q-space samplings (though we have focused on multi-shells). As opposed to the original HYDI-DSI, the quantitative indices derived from these EAPs are relatively robust to the maximum b-value acquired, since the bandwidth of the signal is estimated from a DTI approximation fitted to the low b-value regimen instead of from the sampling protocol itself. Of course, smoothness constraints impose a reduction of the bandwidth of the estimated EAP as the maximum acquired b-value decreases. The keystone of HYDI-DSI-QP is replacing the re-gridding and linear interpolation of the q-space with a constrained, Laplacian-regularized optimization problem, which indeed makes the algorithm more time-efficient.

In fact, the computational complexity in the original HYDI-DSI approach becomes unnecessarily increased by its straightforward implementation with a voxel-wise call to Matlab's `griddatan`. Note the slowest part of this method is the computation of the convex hull and Delaunay's triangulation: if the lattice layout is kept constant with respect to the q-space samples (i.e. it is non-adaptive with regard to the DTI approximation), these two computations could be done once for the entire volume, and the only repeated operation would be the (very fast) linear interpolation. Note, however, this would not apply in case we want to accurately compute directional indices (e.g. RTAP and RTPP) or design adaptive bandwidths like in HYDI-DSI-AB. In the former case, the lattice will become rotated with respect to the q-space samples to align its axes with the principal diffusion directions; in the latter, the lattice will be both rotated and stretched. In both cases, the proposed HYDI-DSI-QP will be a faster option.

An additional advantage of the proposed HYDI-DSI-QP is the possibility it grants to impose positivity constraints, which is not feasible with interpolation-based methods. Fig. 5 highlights this is certainly a non-negligible issue, as long as the greatest impact of negative-valued EAP samples shows up within relevant white matter structures such as the CC or the CST. Moreover, this problem worsens for more modest (hence, closer to what may be found in clinical applications) q-space samplings than the outstanding MGH HCP data set. Remarkably, it has been shown that non-negativity constraints unleash their highest potential when the underlying function representing the EAP is enforced to be non-negative in its entire continuous domain (Dela Haije et al., 2020). With HYDI-DSI-QP, on the contrary, only a discrete subset of the EAP (the lattice nodes) is assured to be non-negative. Note this is inherent to DSI-like, non-parametric approaches, for which the discrete values to be estimated are not parameters to reconstruct a continuous (non-negative) function, but sparse values of the function itself.

The problems of q-space re-gridding have been previously recognized in the Generalized DSI (Tian et al., 2019, GDSI), where the authors work around the problem with a direct discretization of Eq. (1) in the form of a quadrature rule for Fourier's integral. There, each q-space sample is weighted accounting for the volume it occupies, in a way that there is no need for further interpolation. It has the obvious advantage of being computationally very efficient, since it reduces to a (fixed) matrix product at each voxel. Besides, the same scheme easily fits Cartesian or spherical lattices in both the q-space and the EAP domain. However, it is by no means free of certain problems: first, the bandwidth of the signal is directly limited by the maximum b-value acquired, as it is with HYDI-DSI, which in practice means this method will be well suited only for specific data sets like the HCP MGH. Second, the raw discretization of Eq. (1), as opposed to the DFT computed over a q-space signal forced to evaluate to 1 at $\mathbf{q} = \mathbf{0}$, does not guarantee the EAP to have unit mass, which is indeed the case for all the methods explored in this paper. Of course, the positivity of the EAP is neither assured *per se*. Finally, for more general

samplings than Cartesian or spherical, it is not clear how the volume-dependent weights of GDSI would be computed. These samplings would not represent any particular issue for HYDI-DSI-like approaches.

A different approach dealing with q-space interpolation was proposed by Yeh and Verstynen (2016), where the authors drew HARDI data from DSI-like or multi-shell samplings by interpolation. This is based on the linear relation between the so-called Spin Distribution Function (SDF) and the attenuation signal samples. At the very end, the problem is solved as a Tikhonov-regularized, non-constrained least squares problem. The authors claim that non-negativity constraints become unnecessary resorting to empirical considerations, likely owing to the fact that, being aimed at reducing DSI-like/multi-shell data to HARDI schemes, the problem they address is vastly over-determined. In any case, this approach focuses on HARDI analysis techniques, so that it does not stand a direct comparison with our proposal.

The price to pay when getting rid of q-space interpolation is the introduction of two additional algorithm parameters (besides the lattice size) that were not present in the original HYDI-DSI: the bandwidth selection μ and the Laplacian weighting λ . Fortunately, Figs. 3 and 4 show they can be fixed in a *one size fits all* fashion, regardless of the actual q-space sampling used. Indeed, their values remain constant in Figs. 7–10 as the outermost acquired shells are progressively removed, and even so HYDI-DSI-QP exhibits a stable behavior. Moreover, all the results reported throughout the paper used $\mu = 0.05$ and $\lambda = 0.5$ after Figs. 3 and 4, without any further optimization.

Regarding the actual accuracy of HYDI-DSI-like approaches, Fig. 6 shows the three of them provide excellent approximations for the full HCP MGH protocol (with b-values up to $10,000 \text{ s/mm}^2$), unless very poor PSNR values are considered. Though the estimations obtained with either of the three methods strongly correlate with ground-truth values, it remains clear that the proposed HYDI-DSI-QP performs the best at estimating the actual order of magnitude of the RTxP indices, with logarithmic scale-shifts near 0. The actual advantage of our proposal, however, is demonstrated in Figs. 7–10: HYDI-DSI-QP clearly outperforms the original HYDI-DSI in the estimation of RTxP values as fewer shells are available, hence proving itself more useful with non-dedicated, more clinically-suitable diffusion data sets.

Remarkably, the progressive removal of the outermost shells results in two artifacts: first, since fewer samples are available for the estimation, the variance of the estimation (the width of the points clouds in the joint histograms) increases, as expected. Second, the measures are scale-shifted with respect to their values at $b = 10,000 \text{ s/mm}^2$. This is evident with HYDI-DSI, since the calculated bandwidth is directly provided by the maximum b-value acquired. With HYDI-DSI-QP, the reason may be found in Eq. (16): since no high-frequency samples are available, this part of the spectrum is governed by the penalty term, which promotes lower-pass solutions, hence decreasing the values of the RTxP. With HYDI-DSI-AB, on the contrary, the non-measured part of the spectrum will be linearly interpolated, so that the bandwidth can even result artificially increased. Note this is a major difference of our approach when compared to CS-based proposals (Bilgic et al., 2012; Menzel et al., 2011; Young et al., 2017): CS will provide faithful, full-bandwidth approximations to the EAP even if the q-space is sampled below Nyquist's rate (but up to a sufficiently large b-value). This is attained through the use of a *sparsifying* transform and ℓ_1 -based optimization procedures. Our HYDI-DSI-QP avoids the need for such transform by directly computing samples of the function of interest (the discrete EAP), which at the same time allows to directly impose positivity constraints in a straightforward manner.

Note the computations of the RTPP (see Fig. 7), the RTAP (Fig. 8), the RTPP (Fig. 9), and the MSD (Fig. 10) are quite robust to the elimination of large b-valued shells with HYDI-DSI-QP. In other words, HYDI-DSI-QP, as compared to the original HYDI-DSI, grants the opportunity to accurately estimate non-parametric, DSI-like information from non-specific samplings. Yet, HYDI-DSI-QP provides pretty decent

approximations of the bronze standard from even a unique shell at $b = 1000\text{s/mm}^2$, which could allow our proposal to compete with recent model-based/parametric approaches devised for quantitative dMRI analysis from clinically feasible setups (Aja-Fernández et al., 2020, 2021, 2022).

More interestingly, the experiments presented throughout the paper evidence that our HYDI-DSI-QP attains, at the very least, comparable results to those obtained with MAPL in terms of accuracy, robustness, and reliability. Indeed, we are able to clearly and consistently outperform MAPL in many scenarios (compare top and bottom rows of the Bland and Altman graphs in Figs. 7 through 10). Though we cannot claim that MAPL is the state of the art in EAP imaging (to our knowledge, there are not systematic and exhaustive comparative studies supporting the better performance of MAPL compared to other techniques), it is undoubtedly the most popular approach for multi-shells, despite being extremely time consuming and resource demanding. In this sense, we have designed an EAP imaging technique that attains a speedup at least $100\times$ over MAPL with a comparable or even better performance.

Compared to its predecessor (HYDI-DSI by Wu et al. (2008)), HYDI-DSI-QP no longer requires *ad hoc* q-space sampling schemes, but it properly works with those currently available in public databases. Putting all together, HYDI-DSI-QP might constitute a milestone for bringing up the potential of advanced EAP imaging to the analysis of medium-large sized databases.

In this work, we have focused on quantitative dMRI, which is also the main topic covered by Wu et al. (2008), i.e. we aim at accurately computing certain indices derived from diffusion measurements that can potentially reflect features and processes taking place in the white matter at a micro-structure level. As such, we have not paid attention to the computation of ODF fields or EAP-based tractography, which is in turn an important matter of concern in the related literature (Assemlal et al., 2009; Hosseinbor et al., 2013; Tian et al., 2019; Tristán-Vega and Aja-Fernández, 2021). The computation of the EAP in a Cartesian lattice, especially when it is oriented following the principal diffusion directions, nicely fits the computation of the usual indices (RTOP, RTAP, RTPP, MSD), but a spherical arrangement is better suited for ODF or directionality description. Wu and Alexander (2007), devising HYDI as a put-all-together method, were not concerned about this pitfall because ODF/tractography were supposed to be computed independently with one of the available DTI/HARDI techniques. If this complementary information has to be extracted from the Cartesian EAP itself, some sort of re-gridding/interpolation in the EAP domain is required. This might well be a linear interpolation of the EAP samples themselves, or a more sophisticated approach could be thought of: proceeding as in the Appendix, we could obtain a set of Fourier coefficients of the periodically extended EAP, which is equivalent to a re-gridding of the q-space. Then, the analogous developments as in Section 2.3 would allow building a matrix relating the Cartesian sampling of the q-space to the desired spherical grid in the EAP domain, from which computing ODFs would now be trivial. Note that both MAP-MRI and MAPL face similar issues, since they equally rely on functions naturally defined in the Cartesian domain (Fick et al., 2016b; Özarslan et al., 2013b).

Yet, the RTOP, the RTAP, the RTPP and the MSD are not the unique EAP-derived quantities of interest within dMRI: the Non-Gaussianity (NG) and the Propagator Anisotropy (PA), which are naturally derived from the MAP-MRI representation (Özarslan et al., 2013b), have proven themselves certainly valuable for the description of several pathologies (Avram et al., 2016; Bernstein, 2019; Fick et al., 2016a). For example, the PA is defined as the *distance* from the EAP to its equivalent isotropic signal, defined as the spherical average of $P(\mathbf{R})$ at each $\|\mathbf{R}\|$. Once again, the Cartesian description of the EAP with HYDI-DSI approaches is a pitfall in the computation of this measurement that needs to be solved via numeric interpolation. Since HYDI-DSI-QP

pursues the analytical computation of quantitative dMRI indices (see Section 3), the need for numeric interpolation in all these cases is a clear limitation.

The proposed method is aimed at the non-parametric description of the EAP. The relevance of this particular feature arguably relies on the expectation that a non-parametric representation provides extra DoF over parametric ones, meaning that the derived indices will hopefully exhibit a higher sensitivity to micro-structure changes inside the white matter. Conversely, more DoF are likely to make the EAP representation more prone to random drifts induced by noise and imaging artifacts, i.e. reduced specificity. With regard to the latter issue, we have characterized specificity by means of the inter-session coefficient of variation (hither defined as the reproducibility), which should remain relatively low for the derived dMRI measures to be clinically feasible. This criterion has been met for all measures estimated under HYDI-DSI and HYDI-DSI-QP approaches as stated in Table 1(a), with coefficients of variation always below 5%.

However, the analysis in terms of raw reproducibility, despite being relatively simple, can be misleading: Table 1(a) suggests that the original HYDI-DSI might be clearly preferable for the estimation of RTOP and RTAP, since it is twice as repeatable as HYDI-DSI-QP for these indices. However, the previous experiments highlight the fact that the indices estimated with either technique are not directly comparable due to the very different bandwidths estimated for the EAP in each case. In other words, the heavy scale-shifts HYDI-DSI introduces in RTxP values as a consequence of abruptly cropping the bandwidth of the EAP can be related to sensitivity losses: Table 1(b) measures this effect by normalizing inter-session coefficients of variation with inter-subject differences. With this normalization, it remains clear that all the methods compared are equally reliable for the estimation of RTOP and RTPP, meanwhile HYDI-DSI-QP is clearly preferable for RTPP and MSD. In any case, the reproducibility and reliability studies presented in Table 1 demonstrated promising potential of the measures to be further transferred to the clinical domain or to be used in a neurodevelopment research scenario, e.g. brain aging or longitudinal studies.

Finally, the analysis of dMRI measures in terms of reliability (not just reproducibility) is a novel contribution of this paper, which can pave the way for the selection of an appropriate sample size to preserve a trade-off between a long acquisition time or group size and a high significance of statistical tests (Zuo et al., 2019). Such trade-off will be pursued through the choice for an appropriate EAP reconstruction method, either parametric or non-parametric, and, as we have illustrated here, an appropriate design of the bandwidth of the EAP. Besides, it seems likely that it will strongly depend on the particular characteristics of the database to be analyzed. In this sense, a wider study generalizing the one carried out in Section 5.8 would allow the systematic comparison of parametric and non-parametric EAP imaging techniques with regard to their reliability in different situations (i.e. for different imaging protocols). Including HYDI-DSI-like approaches in the pool of compared methods is now possible after the present paper, since we have demonstrated that HYDI-DSI-QP relies notably less on dedicated q-space samplings than the original HYDI-DSI does, so that it can be computed over existing test-retest databases (Koller et al., 2021; Van Essen et al., 2013) as-it-is.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

We have used only publicly available databases. The code is freely available for download.

Acknowledgments

This work was supported by the Ministerio de Ciencia e Innovación of Spain with research grants PID2021-124407NB-I00 and TED2021-130758B-I00.

Tomasz Pieciak acknowledges the Polish National Agency for Academic Exchange for grant PPN/BEK/ 2019/1/00421 under the Bekker programme and the Ministry of Science and Higher Education (Poland) under the scholarship for outstanding young scientists (692/STYP/13/2018).

Data collection and sharing for this project was partially provided by the *Human Connectome Project*, <https://ida.loni.usc.edu/login.jsp> (HCP; Principal Investigators: Bruce Rosen, M.D., Ph.D., Martinos Center at Massachusetts General Hospital, Arthur W. Toga, Ph.D., University of Southern California, Van J. Weeden, MD, Martinos Center at Massachusetts General Hospital). 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 is the result of efforts of co-investigators from the University of Southern California, Martinos Center for Biomedical Imaging at Massachusetts General Hospital (MGH), Washington University, and the University of Minnesota. HCP data are disseminated by the Laboratory of Neuro Imaging at the University of Southern California.

Data were also provided in part by the *Human Connectome Project*, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

Appendix. Description of the Laplacian penalty

The Laplacian of the EAP, as required in Eq. (16), could be approximated with finite differences computed over \mathbf{P} itself. Instead, we look for an *exact* representation based on the DFT. In order to avoid discontinuities due to the inherent periodic boundary conditions, we extend the original EAP lattice to a new one with size $(2N_x + 2) \times (2N_y + 2) \times (2N_z + 2)$ by zero padding at $N_{e_j} + 1$. By rearranging the antipodal symmetric $P_{\Theta}(k/Q_x, l/Q_y, m/Q_z)$ ($\{k, l, m\} = -N_{e_j} \dots N_{e_j} + 1$), we get a column vector \mathbf{P}_{sym} that relates to the signal in the q-space as:

$$\mathbf{P}_{\text{sym}} = Q W_{\text{sym}} \mathbf{E}_{\text{sym}}, \quad (\text{A.1})$$

where the column vector \mathbf{E}_{sym} is rearranged from the signal $E_{\Theta}(Q_x u/2(N_x + 1), Q_y v/2(N_y + 1), Q_z w/2(N_z + 1))$ ($\{u, v, w\} = -N_{e_j} \dots N_{e_j} + 1$), and W_{sym} is the DFT matrix whose entries have the form (Oppenheim et al., 1997):

$$[W_{\text{sym}}]_{r(u,v,w),c(u,v,w)} = \exp\left(-i\pi\left(ku/N_x + 1 + lv/N_y + 1 + mw/N_z + 1\right)\right). \quad (\text{A.2})$$

Since W_{sym} has Hermitian symmetry, its inversion becomes trivial. Besides, since both the EAP and the attenuation signal are necessarily real, we get:

$$\mathbf{E}_{\text{sym}} = \frac{1}{Q} \Re\{W_{\text{sym}}^H\} \mathbf{P}_{\text{sym}}. \quad (\text{A.3})$$

The computation of the Laplacian can be done by relating this operator to its dual in the q-space domain, as it has been thoroughly exploited in the literature (Tristán-Vega et al., 2009, 2010). Hence:

$$\begin{aligned} \mathbf{F}_{\text{sym}} &= -4\pi^2 \text{diag}(\mathbf{s}_{\text{sym}}) \mathbf{E}_{\text{sym}} \\ &= \frac{-4\pi^2}{Q} \text{diag}(\mathbf{s}_{\text{sym}}) \Re\{W_{\text{sym}}^T\} \mathbf{P}_{\text{sym}} = \frac{1}{Q} \Re\{W_f^T\} \mathbf{L}_{\text{sym}}, \end{aligned} \quad (\text{A.4})$$

where \mathbf{L}_{sym} stands for the Laplacian of the EAP, sampled at the same lattice points as the EAP itself, and rearranged as a column vector.

Vector \mathbf{s}_{sym} represents the (column-rearranged) squared modules of the lattice nodes in the q-space, i. e.:

$$[\mathbf{s}_{\text{sym}}]_{r(u,v,w)} = \frac{Q_x^2 u^2}{4(N_x + 1)^2} + \frac{Q_y^2 v^2}{4(N_y + 1)^2} + \frac{Q_z^2 w^2}{4(N_z + 1)^2}. \quad (\text{A.5})$$

Therefore, \mathbf{F}_{sym} is the DFT of the signal corresponding to the Laplacian of the sampled EAP. According to Parseval's theorem, the energy of the former equals the energy of the latter. By identifying terms in Eqs. (16) and (A.4) we can derive:

$$\mathcal{L}_{\text{sym}} = \frac{-4\pi^2}{Q} \text{diag}(\mathbf{s}_{\text{sym}}) \Re\{W_{\text{sym}}^T\}. \quad (\text{A.6})$$

It only remains to drop off the rows and columns of W_{sym}^T corresponding to antipodal symmetries (like we did in Section 2.3), as well as those columns corresponding to the original zero-padding of the EAP at extra lattice points, to get W^T (and perform analogous operations with the rows of \mathbf{s}_{sym} to get \mathbf{s}). Finally:

$$\mathcal{L} = Q^{-2/3} \frac{-4\pi^2}{Q} \text{diag}(\mathbf{s}) \Re\{W^T\} = -4\pi^2 Q^{-5/3} \text{diag}(\mathbf{s}) \Re\{W^T\}, \quad (\text{A.7})$$

where the additional normalization factor $Q^{-2/3}$ obeys to the need for dimensional homogeneity between \mathcal{F} and \mathcal{L} in Eq. (16).

References

Aja-Fernández, S., de Luis-García, R., Afzali, M., Molendowska, M., Pieciak, T., Tristán-Vega, A., 2020. Micro-structure diffusion scalar measures from reduced MRI acquisitions. *PLoS One* 15 (3), 1–25.

Aja-Fernández, S., Pieciak, T., Martín-Martín, C., Á, Planchuelo-Gómez, de Luis-García, R., Tristán-Vega, A., 2022. Moment-based representation of the diffusion inside the brain from reduced DMRI acquisitions: Generalized AMURA. *Med. Image Anal.* 77, 102356.

Aja-Fernández, S., Tristán-Vega, A., Jones, D.K., 2021. Apparent propagator anisotropy from single-shell diffusion MRI acquisitions. *Magn. Reson. Med.* 85 (5), 2869–2881.

Assaf, Y., Basser, P.J., 2005. Composite hindered and restricted model of diffusion (CHARMED) MR imaging of the human brain. *NeuroImage* 27 (1), 48–58.

Assemlal, H.-E., Tschumperlé, D., Brun, L., 2009. Efficient and robust computation of PDF features from diffusion MR signal. *Med. Image Anal.* 13 (5), 715–729.

Avram, A.V., Sarlls, J.E., Barnett, A.S., Özarslan, E., Thomas, C., Irfanoglu, M.O., Hutchinson, E., Pierpaoli, C., Basser, P.J., 2016. Clinical feasibility of using mean apparent propagator (MAP) MRI to characterize brain tissue microstructure. *NeuroImage* 127, 422–434.

Barber, C.B., Dopkin, D.P., Huhdanpaa, H.T., 1996. The quickhull algorithm for convex hulls. *ACM Trans. Math. Software* 22 (4), 469–483.

Basser, P.J., Mattiello, J., Le Bihan, D., 1994. Estimation of the effective self-diffusion tensor from the NMR spin echo. *J. Magn. Reson.* 103 (3), 247–254.

Basser, P.J., Pierpaoli, C., 1996. Microstructural and physiological features of tissues elucidated by Quantitative-Diffusion-Tensor MRI. *J. Magn. Reson.* 111 (3), 209–219.

Benjamini, D., Basser, P.J., 2019. Water mobility spectral imaging of the spinal cord: Parametrization of model-free Laplace MRI. *Magn. Reson. Imaging* 56, 187–193.

Bernstein, A.S., 2019. *Advanced Diffusion MRI Techniques: Methodological Development and Clinical Application* (Ph.D. thesis). The University of Arizona, Tucson (AZ), United States of America.

Bilgic, B., Setsompop, K., Cohen-Adad, J., Yendiki, A., Wald, L.L., Adalsteinsson, E., 2012. Accelerated diffusion spectrum imaging with compressed sensing using adaptive dictionaries. *Magn. Reson. Med.* 68 (6), 1747–1754.

Bland, J.M., Altman, D.G., 1986. Statistical methods for assessing agreement between two methods of clinical measurement. *Lancet* 1, 307–310.

Boscolo Galazzo, I., Brusini, L., Obertino, S., Zucchelli, M., Granziera, C., Menegaz, G., 2018. On the viability of diffusion MRI-based microstructural biomarkers in ischemic stroke. *Front. Neurosci.* 12 (92).

Brusini, L., Obertino, S., Galazzo, I.B., Zucchelli, M., Krueger, G., Granziera, C., Menegaz, G., 2016. Ensemble average propagator-based detection of microstructural alterations after stroke. *Int. J. Comput. Assist. Radiol. Surg.* 11 (9), 1585–1597.

Callaghan, P., 1991. *Principles of Nuclear Magnetic Resonance Microscopy*. Clarendon Press, Oxford.

Caruyer, E., Deriche, R., 2012. Diffusion MRI signal reconstruction with continuity constraint and optimal regularization. *Med. Image Anal.* 16 (6), 1113–1120.

Daianu, M., Jacobs, R.E., Weitz, T.M., Town, T.C., Thompson, P.M., 2015. Multi-shell hybrid diffusion imaging (HYDI) at 7 Tesla in TgF344-AD transgenic Alzheimer rats. *PLoS One* 10 (12), e0145205.

Dela Haije, T., Özarslan, E., Feragen, A., 2020. Enforcing necessary non-negativity constraints for common diffusion MRI models using sum of squares programming. *NeuroImage* 209, 116405.

- Descoteaux, M., Angelino, E., Fitzgibbons, S., Deriche, R., 2007. Regularized, fast, and robust analytical Q-Ball imaging. *Magn. Reson. Med.* 58, 497–510.
- Descoteaux, M., Deriche, R., Le Bihan, D., Mangin, J.-F., Poupon, C., 2011. Multiple q-shell diffusion propagator imaging. *Med. Image Anal.* 15 (4), 603–621.
- Donoho, D.L., 2006. Compressed sensing. *IEEE Trans. Inform. Theory* 52 (4), 1289–1306.
- Fan, Q., Witzel, T., Nummenmaa, A., Van Dijk, K.R., Van Horn, J.D., Drews, M.K., Somerville, L.H., Sheridan, M.A., Santillana, R.M., Snyder, J., Hedden, T., Shaw, E.E., Hollinshead, M.O., Renvall, V., Zanzonico, R., Keil, B., Cauley, S., Polimeni, J.R., Tisdall, D., Buckner, R.L., Wedeen, V.J., Wald, L.L., Toga, A.W., Rosen, B.R., 2016. MGH-USC human connectome project datasets with ultra-high b-value diffusion MRI. *NeuroImage* 124, 1108–1114.
- Fick, R.H., Daiuan, M., Pizzolato, M., Wassermann, D., Jacobs, R.E., Thompson, P.M., Town, T., Deriche, R., 2016a. Comparison of biomarkers in transgenic Alzheimer rats using multi-shell diffusion MRI. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, Athens, Greece, pp. 187–199.
- Fick, R.H., Wassermann, D., Caruyer, E., Deriche, R., 2016b. MAPL: Tissue microstructure estimation using Laplacian-regularized MAP-MRI and its application to HCP data. *NeuroImage* 134, 365–385.
- Froeling, M., Tax, C.M., Vos, S.B., Luijten, P.R., Leemans, A., 2017. MASSIVE brain dataset: Multiple acquisitions for standardization of structural imaging validation and evaluation. *Magn. Reson. Med.* 77 (5), 1797–1809.
- Hosseinbor, A.P., Chung, M.K., Wu, Y.-C., Alexander, A.L., 2013. Bessel Fourier orientation reconstruction (BFOR): An analytical diffusion propagator reconstruction for hybrid diffusion imaging and computation of q-space indices. *NeuroImage* 64, 650–670.
- Hosseinbor, A.P., Chung, M.K., Wu, Y.-C., Fleming, J.O., Field, A.S., Alexander, A.L., 2012. Extracting quantitative measures from EAP: A small clinical study using BFOR. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, Nice, France, pp. 280–287.
- Jack, Jr., C.R., Bernstein, M.A., Fox, N.C., Thompson, P., Alexander, G., Harvey, D., Borowski, B., Britson, P.J., L. Whitwell, J., Ward, C., Dale, A.M., Felmlee, J.P., Gunter, J.L., Hill, D.L., Killiany, R., Schuff, N., Fox-Bosetti, S., Lin, C., Studholme, C., DeCarli, C.S., Krueger, G., Ward, H.A., Metzger, G.J., Scott, K.T., Malozzi, R., Blezek, D., Levy, J., Debbins, J.P., Fleisher, A.S., Albert, M., Green, R., Bartzokis, G., Glover, G., Mugler, J., Weiner, M.W., 2008. The Alzheimer's disease neuroimaging initiative (ADNI): MRI methods. *J. Magn. Reson. Imaging* 27 (4), 685–691.
- 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.
- Jenkinson, M., Smith, S., 2001. A global optimisation method for robust affine registration of brain images. *Med. Image Anal.* 5 (2), 143–156.
- Jeurissen, B., Tournier, J.-D., Dhollander, T., Connelly, A., Sijbers, J., 2014. Multi-tissue constrained spherical deconvolution for improved analysis of multi-shell diffusion MRI data. *NeuroImage* 103, 411–426.
- Kaden, E., Kelm, N.D., Carson, R.P., Does, M.D., Alexander, D.C., 2016. Multi-compartment microscopic diffusion imaging. *NeuroImage* 139, 346–359.
- Koller, K., Rudrapatna, U., Chamberland, M., Raven, E.P., Parker, G.D., Tax, C.M., Drakesmith, M., Fasano, F., Owen, D., Hughes, G., Charron, C., Evans, C.J., Jones, D.K., 2021. MICRA: Microstructural image compilation with repeated acquisitions. *NeuroImage* 225, 117406.
- Lampinen, B., Szczepankiewicz, F., Moartensson, J., van Westen, D., Hansson, O., Westin, C.-F., Nilsson, M., 2020. Towards unconstrained compartment modeling in white matter using diffusion-relaxation MRI with tensor-valued diffusion encoding. *Magn. Reson. Med.* 84 (3), 1605–1623.
- Luenberger, D.G., Ye, Y., 2008. Linear and nonlinear programming. In: *International Series on Operations Research & Management Science*, third ed. Springer.
- Menzel, M.I., Tan, E.T., Khare, K., Sperl, J.I., King, K.F., Tao, X., Hardy, C.J., Marinelli, L., 2011. Accelerated diffusion spectrum imaging in the human brain using compressed sensing. *Magn. Reson. Med.* 66 (5), 1226–1233.
- Merlet, S.L., Deriche, R., 2013. Continuous diffusion signal, EAP and ODF estimation via compressive sensing in diffusion MRI. *Med. Image Anal.* 17 (5), 556–572.
- Mori, S., Wakana, S., Van Zijl, P.C., Nagae-Poetscher, L., 2005. *MRI Atlas of Human White Matter*. Elsevier.
- Muller, J., Middleton, D., Alizadeh, M., Zabrecky, G., Wintering, N., Bazzan, A.J., Lang, J., Wu, C., Monti, D.A., Wu, Q., Newberg, A.B., Mohamed, F.B., 2021. Hybrid diffusion imaging reveals altered white matter tract integrity and associations with symptoms and cognitive dysfunction in chronic traumatic brain injury. *NeuroImage: Clin.* 30, 102681.
- Ning, L., Westin, C.-F., Rathi, Y., 2015. Estimating diffusion propagator and its moments using directional radial basis functions. *IEEE Trans. Med. Imag.* 34 (10), 2058–2078.
- Novikov, D.S., Kiselev, V.G., Jespersen, S.N., 2018. On modeling. *Magn. Reson. Med.* 79 (6), 3172–3193.
- Oppenheim, A.V., Willsky, A.S., Nawab, S.H., 1997. *Signals and systems*. In: *Signal Processing Series*, second ed. Prentice Hall.
- Özarslan, E., Koay, C.G., Basser, P.J., 2013a. Simple harmonic oscillator based reconstruction and estimation for one-dimensional q-space magnetic resonance (1D-SHORE). In: *Excursions in Harmonic Analysis*. Vol. 2. Springer, pp. 373–399.
- Özarslan, E., Koay, C.G., Shepherd, T.M., Komlosh, M.E., Irfanoğlu, M.O., Pierpaoli, C., Basser, P.J., 2013b. Mean apparent propagator (MAP) MRI: a novel diffusion imaging method for mapping tissue microstructure. *NeuroImage* 78, 16–32.
- Pfeuffer, J., Provencher, S.W., Gruetter, R., 1999. Water diffusion in rat brain in vivo as detected at very large b-values is multicompartmental, magnetic resonance materials in physics. *Biol. Med.* 8 (2), 98–108.
- Planchuelo-Gómez, Á., García-Azorín, D., Guerrero, Á.L., Aja-Fernández, S., Rodríguez, M., de Luis-García, R., 2021. Multimodal fusion analysis of structural connectivity and gray matter morphology in migraine. *Hum. Brain Mapp.* 42 (4), 908–921.
- Planchuelo-Gómez, Á., García-Azorín, D., Guerrero, Á.L., de Luis-García, R., Rodríguez, M., Aja-Fernández, S., 2020. Alternative microstructural measures to complement diffusion tensor imaging in migraine studies with standard MRI acquisition. *Brain Sci.* 10 (10), 711.
- Salvador, R., Peña, A., Menon, D.-K., Carpenter, T.-A., Pickard, J.-D., Bullmore, E.-T., 2005. Formal characterization and extension of the linearized diffusion tensor model. *Hum. Brain Mapp.* 24, 144–155.
- Smith, S.M., Jenkinson, M., Woolrich, M.W., Beckmann, C.F., Behrens, T.E., Johansen-Berg, H., Bannister, P.R., De Luca, M., Drobnjak, I., Flitney, D.E., et al., 2004. Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage* 23, S208–S219.
- Stejskal, E.-O., Tanner, J.-E., 1965. Spin diffusion measurements: Spin echoes in the presence of a time-dependent field gradient. *J. Chem. Phys.* 42, 288–292.
- Tax, C.M., Grussu, F., Kaden, E., Ning, L., Rudrapatna, U., Evans, C.J., St-Jean, S., Leemans, A., Koppers, S., Merhof, D., et al., 2019. Cross-scanner and cross-protocol diffusion MRI data harmonisation: A benchmark database and evaluation of algorithms. *NeuroImage* 195, 285–299.
- Tian, Q., Yang, G., Leuze, C., Rokem, A., Edlow, B.L., McNab, J.A., 2019. Generalized diffusion spectrum magnetic resonance imaging (GDSI) for model-free reconstruction of the ensemble average propagator. *NeuroImage* 189, 497–515.
- Topgaard, D., 2017. Multidimensional diffusion MRI. *J. Magn. Reson.* 275, 98–113.
- Tristán-Vega, A., Aja-Fernández, S., 2021. Efficient and accurate EAP imaging from multi-shell dMRI with micro-structure adaptive convolution kernels and dual Fourier integral transforms (MiSFIT). *NeuroImage* 227, 117616.
- Tristán-Vega, A., Westin, C.-F., Aja-Fernández, S., 2009. Estimation of fiber orientation probability density functions in high angular resolution diffusion imaging. *NeuroImage* 47 (2), 638–650.
- Tristán-Vega, A., Westin, C.-F., Aja-Fernández, S., 2010. A new methodology for the estimation of fiber populations in the white matter of the brain with the Funk-Radon transform. *NeuroImage* 49 (2), 1301–1315.
- Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Ugurbil, K., 2013. The WU-Minn human connectome project: An overview. *NeuroImage* 80, 62–79.
- Vemuri, B.C., Sun, J., Banerjee, M., Pan, Z., Turner, S.M., Fuller, D.D., Forder, J.R., Entezari, A., 2019. A geometric framework for ensemble average propagator reconstruction from diffusion MRI. *Med. Image Anal.* 57, 89–105.
- Watson, D.F., 1994. *Contouring: a Guide to the Analysis and Display of Spatial Data*. Pergamon.
- Wedeen, V., Hagmann, P., Tseng, W.-Y., Reese, T., Weisskoff, R., 2005. Mapping complex tissue architecture with diffusion spectrum imaging. *Magn. Reson. Med.* 54, 1377–1386.
- Wu, Y.-C., Alexander, A.L., 2007. Hybrid diffusion imaging. *NeuroImage* 36 (3), 617–629.
- Wu, Y., Field, A.S., Alexander, A.L., 2008. Computation of diffusion function measures in q-space using magnetic resonance hybrid diffusion imaging. *IEEE Trans. Med. Imaging* 27 (6), 858–865.
- Wu, Y.-C., Field, A.S., Whalen, P.J., Alexander, A.L., 2011. Age- and gender-related changes in the normal human brain using hybrid diffusion imaging (HYDI). *NeuroImage* 54 (3), 1840–1853.
- Wu, Y.-C., Mustafi, S.M., Harezlak, J., Kodiweera, C., Flashman, L.A., McAllister, T.W., 2018. Hybrid diffusion imaging in mild traumatic brain injury. *J. Neurotrauma* 35 (20), 2377–2390.
- Yeh, F.-C., Verstyne, T.D., 2016. Converting multi-shell and diffusion spectrum imaging to high angular resolution diffusion imaging. *Front. Neurosci.* 10 (418).
- Young, R.J., Tan, E.T., Peck, K.K., Jenabi, M., Karimi, S., Brennan, N., Rubel, J., Lyo, J., Shi, W., Zhang, Z., Prastawa, M., Liu, X., Sperl, J.I., Fatovic, R., Marinelli, L., Holodny, A.I., 2017. Comparison of compressed sensing diffusion spectrum imaging and diffusion tensor imaging in patients with intracranial masses. *Magn. Reson. Imaging* 36, 24–31.
- 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.
- Ziegler, E., Rouillard, M., André, E., Coolen, T., Stender, J., Balteau, E., Phillips, C., Garraux, G., 2014. Mapping track density changes in nigrostriatal and extranigral pathways in Parkinson's disease. *Neuroimage* 99, 498–508.
- Zuo, X.-N., Xu, T., Milham, M.P., 2019. Harnessing reliability for neuroscience research. *Nat. Hum. Behav.* 3, 768–771.