

Learning compact q -space representations for multi-shell diffusion-weighted MRI

Daan Christiaens, Lucilio Cordero-Grande, Jana Hutter, Anthony N. Price, Maria Deprez, Joseph V. Hajnal, J-Donald Tournier

Abstract—Diffusion-weighted MRI measures the direction and scale of the local diffusion process in every voxel through its spectrum in q -space, typically acquired in one or more shells. Recent developments in microstructure imaging and multi-tissue decomposition have sparked renewed attention in the radial b -value dependence of the signal. Applications in motion correction and outlier rejection therefore require a compact linear signal representation that extends over the radial as well as angular domain. Here, we introduce SHARD, a data-driven representation of the q -space signal based on spherical harmonics and a radial decomposition into orthonormal components. This representation provides a complete, orthogonal signal basis, tailored to the spherical geometry of q -space and calibrated to the data at hand. We demonstrate that the rank-reduced decomposition outperforms model-based alternatives in human brain data, whilst faithfully capturing the micro- and meso-structural information in the signal. Furthermore, we validate the potential of joint radial-spherical as compared to single-shell representations. As such, SHARD is optimally suited for applications that require low-rank signal predictions, such as motion correction and outlier rejection. Finally, we illustrate its application for the latter using outlier robust regression.

Index Terms—Diffusion-weighted imaging, Multi-shell HARDI, Blind source separation, Dimensionality reduction

I. INTRODUCTION

Diffusion-weighted magnetic resonance imaging (dMRI) is a noninvasive imaging technique that can probe properties of the cellular microstructure in tissue, such as neurite fibre directions in brain white matter. Its underlying principle is based on sensitizing the MRI signal to the Brownian motion of proton spins along a certain direction and scale encoded by the diffusion gradient [1]. When sampled across a dense set of narrow gradient pulses, the resulting dMRI signal measures the 3-D Fourier spectrum of the ensemble average propagator of the diffusion process in each voxel. This spectral domain is known as q -space, in analogy with the k -space concept in conventional MRI [2].

Manuscript received 2 May 2018; revised 30 July 2018; accepted 24 September 2018. Please see [10.1109/TMI.2018.2873736](https://doi.org/10.1109/TMI.2018.2873736) for citation information.

All authors are with the Centre for the Developing Brain, School of Biomedical Engineering & Imaging Sciences, King's College London, King's Health Partners, St. Thomas' Hospital, London SE1 7EH, United Kingdom (e-mail: daan.christiaens@kcl.ac.uk).

This research was funded by the ERC developing Human Connectome Project [Grant Agreement no. 319456] and MRC strategic funds [MR/K006355/1]. This work was additionally supported by the Wellcome/EPSCRC Centre for Medical Engineering at King's College London [WT 203148/Z/16/Z] and by the National Institute for Health Research (NIHR) Biomedical Research Centre at Guy's and St Thomas' NHS Foundation Trust and King's College London. The views expressed are those of the authors and not necessarily those of the NHS, the NIHR or the Department of Health.

Most dMRI data are nowadays acquired with dense sampling over one or more shells in q -space, with radius determined by the b -value of each shell (assuming fixed diffusion time) [1]. Such single- and multi-shell protocols facilitate a straightforward trade-off between high angular and radial resolution; the former being advantageous for tractography, the latter for microstructure imaging. Moreover, the signal in each shell can be efficiently represented in the spherical harmonic (SH) basis [3], which linearizes spherical deconvolution and related techniques [4]–[9]. Recent developments in microstructure modelling [10]–[20] and multi-tissue decomposition [21]–[23] have sparked renewed attention in the radial b -value dependence of the signal. There is therefore a need for compact, linear representations of the dMRI signal that extend over the radial as well as angular domain.

Applications such as motion and distortion correction [24], [25] and outlier rejection [26]–[28] require such signal representations for generating rank-reduced predictions to which the input data can be registered or compared. Variational methods for (super-resolution) image reconstruction [29] can also benefit from linear and compact representations as these simplify the required numeric optimization. Furthermore, compact signal representations may prove useful for image denoising [30] and for multi-site data harmonization [31]. Crucially, signal representations in such applications should be free of biophysical assumptions, as not to affect subsequent analysis and enable comparison of different modelling approaches.

Here, we introduce a model-free representation of the q -space signal based on the spherical harmonics basis within each shell and a linear, orthonormal decomposition in the radial domain. The combined *spherical harmonics and radial decomposition* (SHARD) offers a bespoke signal representation across all shells, hence building a data-driven basis for the q -space signal in the images at hand.

This approach is similar in spirit to other blind source separation methods in dMRI, such as sparse or convex nonnegative spherical factorization [23], [32], [33]. However, while nonnegativity constraints of the factorized orientation distribution functions (ODF) in those techniques are motivated on biophysical grounds, they also—by necessity—lead to increased residuals on the signal representation that may still contain relevant structure. This work does not impose any constraints apart from orthogonality, sacrificing direct biological interpretability in favour of a model-free, unconstrained basis that is better suited for low-rank signal approximations needed in motion correction and related applications, and which forms an effective basis for subsequent, more biologically inspired

analyses.

In contrast to the harmonic oscillator basis [34], [35] and other similar representations [36]–[42], the SHARD basis does not depend on scaling parameters and is implicitly calibrated to the data at hand. As opposed to the cumulant expansion [43], higher-order tensor representations [44]–[46], or Gaussian processes [47], our decomposition is directly compatible with the SH basis and hence inherits its mathematical properties, including its ability to linearize spherical deconvolution.

In this paper, we outline the theoretical underpinnings for the SHARD decomposition and evaluate its representation accuracy in comparison to alternative bases in human brain data with shells spanning across a range of $0 \leq b \leq 10 \text{ ms}/\mu\text{m}^2$. Additionally, we illustrate its use in an example practical application for outlier rejection.

II. THEORY

For fixed diffusion time τ , the diffusion-weighted signal $S(\mathbf{q})$ at encoding $\mathbf{q} \in \mathbb{R}^3$ can be represented in a linear basis as

$$S(\mathbf{q}) = \sum_{j=1}^{\infty} c_j \Psi_j(\mathbf{q}) \quad (1)$$

In this work, we seek a complete, orthonormal basis $\Psi_j(\mathbf{q})$ suitable for unconstrained, low-rank representations of the dMRI signal. In addition, we require the basis functions to be separable in spherical coordinates, and adopt the basis of real symmetric spherical harmonics (SH) $Y_\ell^m(\theta, \phi)$ for the angular part:

$$\Psi_j(\mathbf{q}) = R_\ell^n(q) Y_\ell^m(\hat{\mathbf{q}}) \quad (2)$$

where $q = \|\mathbf{q}\|$, $\hat{\mathbf{q}} = \mathbf{q}/q$, and index j is determined by even SH order ℓ , phase $m \in [-\ell, \ell]$, and radial basis degree n . Crucially, the basis functions $R_\ell^n(q)$ for the radial domain depend on the SH order ℓ but not on the phase m . This ensures that the radial basis can vary between independent SH frequency bands, while remaining invariant to the orientation of the dMRI signal.

Related work has presented different choices of $R_\ell^n(q)$. Hosseinbor et al. [41] used the spherical Bessel functions $j_\ell(\alpha q)$ at frequency α , which arise naturally when casting the q -space inverse Fourier transform in spherical coordinates. However, discretizing the frequency α requires imposing zero-value or zero-derivative Sturm-Liouville boundary conditions at q_{\max} [48], which are generally incompatible with the nature of the dMRI signal and can thus cause aliasing effects. Özarıslan et al. [34] introduced the harmonic oscillator (SHORE) basis of spherical Laguerre polynomials as a higher-order generalization of the diffusion tensor model, with basis functions depending on a scale factor ζ that needs to be tuned to the voxel or image. Finally, many current multi-shell analysis methods that treat shells independently can be regarded as using Dirac-delta basis functions $R_\ell^n(q) = \delta(q - q_n)$ at discrete shell positions q_n [19], [21], [23].

Here we propose to learn an orthonormal function basis for the radial domain from the data, using the singular value decomposition (SVD) of the signal across different shells and in different SH frequency bands ℓ . Because the SH basis functions

Y_ℓ^m are mutually orthogonal between different orders ℓ , it suffices to ensure orthogonality between basis functions R_ℓ^n of different n for fixed ℓ . Hence, we can independently resolve SVD responses in each band ℓ . To this end, we arrange the order- ℓ SH coefficients of all voxels in the image, mask or patch, in a $N_{\text{shells}} \times (N_{\text{vox}} \cdot (2\ell + 1))$ matrix

$$S_\ell = \begin{bmatrix} S_\ell^{(1)} & S_\ell^{(2)} & \dots & S_\ell^{(N_{\text{vox}})} \end{bmatrix} \quad (3)$$

with

$$S_\ell^{(v)} = \begin{bmatrix} s_{\ell, q_1, v}^{-\ell} & \dots & s_{\ell, q_1, v}^m & \dots & s_{\ell, q_1, v}^\ell \\ \vdots & & \vdots & & \vdots \\ s_{\ell, q_N, v}^{-\ell} & \dots & s_{\ell, q_N, v}^m & \dots & s_{\ell, q_N, v}^\ell \end{bmatrix}, \quad (4)$$

where $s_{\ell, q, v}^m = \langle Y_\ell^m(\theta, \phi), S_v(q, \theta, \phi) \rangle$ is the projection of the signal onto the SH basis, i.e., the ℓ, m SH coefficient of shell q in voxel v . Thus, shells q are laid out one per row, voxels v and phase terms m are laid out in columns. The truncated SVD decomposes this matrix

$$S_\ell = \sum_{n=1}^{N_{\text{shells}}} \mathbf{u}_{\ell, n} \sigma_{\ell, n} \mathbf{v}_{\ell, n}^\top = \mathbf{U}_\ell \Sigma_\ell \mathbf{V}_\ell^\top \quad (5)$$

into left and right singular vectors $\mathbf{u}_{\ell, n}$ and $\mathbf{v}_{\ell, n}$, and associated singular values $\sigma_{\ell, n}$ in decreasing order. The vectors $\mathbf{u}_{\ell, n}$ (eigenvectors of $S_\ell S_\ell^\top$) span an orthonormal basis for the order- ℓ SH band across shells, that uniquely determine the rotation-invariant principal components in the signal. The vectors $\mathbf{v}_{\ell, n}$ contain the associated voxel coefficients of each basis component. The singular values $\sigma_{\ell, n}$ measure the effect size of each component.

This decomposition directly extends the isotropic ($\ell = 0$) data representation used in Tournier et al. [49] for optimizing the diffusion gradient sampling scheme to higher SH order $\ell \geq 0$, thus facilitating a generalized signal representation including all angular information. The proposed basis also extends the rotation-invariant signal features that have recently been introduced for microstructural modelling [18]–[20]. In a single voxel, the widely accepted spherical convolution model [7], [20], [50] assumes that the signal can be factorized into an axially symmetric response function $h_{\ell, q}$ (the microstructure) and an orientation distribution function p_ℓ^m (ODF, the mesostructure): $s_{\ell, q}^m = h_{\ell, q} p_\ell^m$. Under this identity, the voxel-wise covariance matrix

$$\begin{aligned} S_\ell^{(v)} S_\ell^{(v)\top} &= \begin{bmatrix} \sum_m s_{\ell, q_1}^m{}^2 & \dots & \sum_m s_{\ell, q_1}^m s_{\ell, q_N}^m \\ \vdots & \ddots & \vdots \\ \sum_m s_{\ell, q_N}^m s_{\ell, q_1}^m & \dots & \sum_m s_{\ell, q_N}^m{}^2 \end{bmatrix} \quad (6) \\ &= \begin{bmatrix} h_{\ell, q_1}^2 & \dots & h_{\ell, q_1} h_{\ell, q_N} \\ \vdots & \ddots & \vdots \\ h_{\ell, q_N} h_{\ell, q_1} & \dots & h_{\ell, q_N}^2 \end{bmatrix} \underbrace{\left(\sum_m p_\ell^m{}^2 \right)}_{p_\ell} \quad (7) \\ &= p_\ell \mathbf{h}_\ell \mathbf{h}_\ell^\top \quad \text{with} \quad \mathbf{h}_\ell = [h_{\ell, q_1} \dots h_{\ell, q_N}]^\top \quad (8) \end{aligned}$$

becomes a rank-1 matrix (in the absence of noise) scaled with the ODF power p_ℓ . The diagonal elements of this matrix [51] or ratios thereof [19] have previously been used as rotation-invariant signal features. Here, we use the full covariance matrix $S_\ell^{(v)} S_\ell^{(v)\top}$, whose principal eigenvector $\mathbf{h}_\ell / \|\mathbf{h}_\ell\|$ captures the




Fig. 1. Covariance matrices $\tilde{S}_\ell \tilde{S}_\ell^\top$ in simulated signals. (a) Simulated free water signal with isotropic diffusivity $D_f = 3.0 \mu\text{m}^2/\text{ms}$. (b) Simulated bi-exponential signal with 50% restricted intra-axonal diffusion with $D_i^\parallel = 2.2 \mu\text{m}^2/\text{ms}$ and 50% hindered extra-axonal diffusion with $D_e^\parallel = 1.2 \mu\text{m}^2/\text{ms}$ and $D_e^\perp = 0.7 \mu\text{m}^2/\text{ms}$. (c) Simulated effect of Gaussian noise on the q -space sampling scheme of dataset 1 after rescaling with the number of samples, reducing to the identity matrix in all overdetermined shells.

microstructural information in the signal up to a scaling factor. Example covariance matrices for (isotropic) free water and bi-exponential white matter are shown in Fig. 1a-b, simulated directly in spherical harmonics using the forward relation for Gaussian signal response in each compartment [6].

When extended across all voxels, the full covariance matrix $S_\ell S_\ell^\top = \sum_v S_\ell^{(v)} S_\ell^{(v)\top}$ is no longer rank-1 and its eigenvectors capture the spatial variation of $h_{\ell,q}$. Therefore, the basis $\mathcal{U} = \{U_\ell | \forall \ell \geq 0\}$ for the radial domain spans the complete microstructural properties encoded in the signal for a given multi-shell protocol. Example matrices $S_\ell S_\ell^\top$ and derived basis functions and effect sizes are illustrated in Fig. 2.

In the remainder of this work, we define basis functions $R_\ell^n(q)$ for the radial domain as discrete functions at shell locations q_i , i.e.,

$$R_\ell^n(q) = \mathbf{u}_{\ell,n} \cdot [\delta(q - q_1), \dots, \delta(q - q_N)]^\top \quad (9)$$

The functions of smallest degree $n \leq N_{\text{shells}}$ capture the strongest covariance between individual shells across all voxels in the image. Note that in general, one may define $R_\ell^n(q)$ using any desired numerical interpolation method between shells.

By virtue of the Eckart-Young SVD theorem, this basis provides the optimal low-rank representation of multi-shell dMRI data: a signal representation in a truncated basis of rank r has minimum Frobenius error w.r.t. the full-rank data. To this end, we can define index $j(\ell, m, n)$ in such a way that all basis functions are sorted in order of decreasing effect size $\varepsilon_j = \frac{\sigma_{\ell,n}^2}{2\ell+1}$, where the normalisation factor accounts for the number of basis functions in each harmonic band ℓ . As such,

the rank- r approximation

$$\hat{S}(\mathbf{q}) = \sum_{j=1}^r c_j \Psi_j(\mathbf{q}) \quad (10)$$

represents the multi-shell dMRI signal with minimal L_2 -error $\|S(\mathbf{q}) - \hat{S}(\mathbf{q})\|_{L_2}$. This property facilitates applications for denoising, motion correction and outlier rejection.

III. MATERIALS AND METHODS

A. Data and preprocessing

1) *Dataset 1*: Multi-shell data were acquired in a healthy subject on a 3 T Philips Achieva TX system using a 32-channel head coil and image-based shimming. The dMRI data were sampled over 11 shells ranging from $b = 0 \text{ s/mm}^2$ to $b = 10,000 \text{ s/mm}^2$, spaced equidistantly in q (note that we use $q \propto \sqrt{b}$ interchangeably throughout this paper). The number of samples in each shell ranges from 15 to 115 respectively, scaled linearly with the shell surface $\propto q^2 \propto b$, resulting in a total of 550 dMRI volumes interleaved for optimal duty cycle effects [52]. Each volume was acquired with single-shot echo planar imaging (EPI), $G_{\text{max}} = 80 \text{ mT/m}$, $\text{TE} = 99 \text{ ms}$, $\text{TR} = 6900 \text{ ms}$, multiband factor $\text{MB} = 2$, SENSE factor 2, isotropic resolution 2.5 mm. In addition, $b = 0$ single-band reference scans were acquired with 4 phase encoding directions, for image reconstruction [53] and for susceptibility-induced distortion correction.

2) *Dataset 2*: A second healthy volunteer was scanned on the same system with higher spatial resolution and more standard b -value range. EPI acquisition parameters are: $\text{MB} = 3$, $\text{SENSE} = 2$, isotropic resolution 2.0 mm, $\text{TE} = 105 \text{ ms}$ and $\text{TR} = 6000 \text{ ms}$. The dMRI gradient table was designed

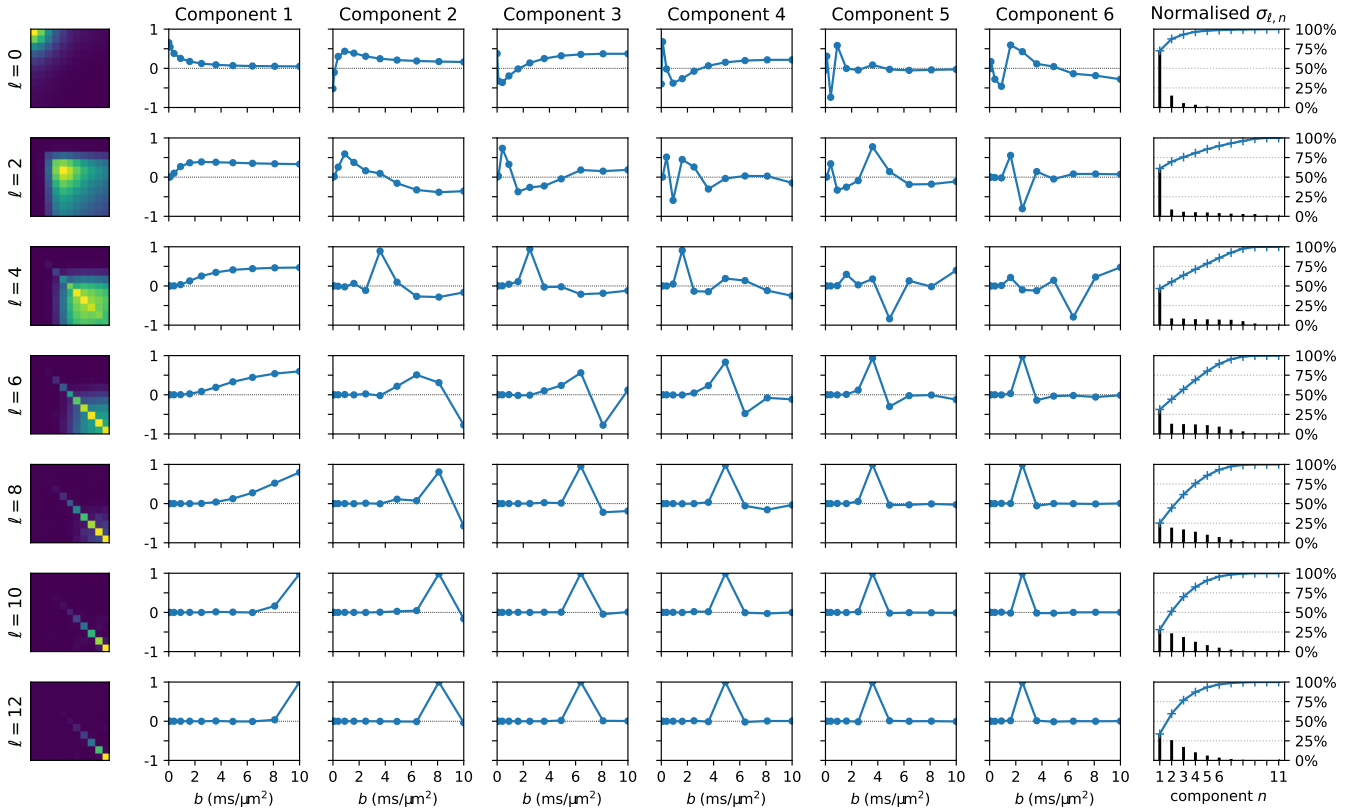


Fig. 2. SHARD decomposition of dataset 1. Rows correspond to spherical harmonic bands of order $\ell = 0, 2, \dots, 12$. The left column illustrates square, symmetric matrices $\tilde{S}_\ell \tilde{S}_\ell^T$, showing the covariance between shells in each harmonic band. The middle columns plot the decomposition basis vectors of the 6 leading components, corresponding to the principal eigenvectors of the matrices on the left. The plots in the right column depict the singular values of each component in decreasing order and their cumulative sum.

with shells at $b = 0, 600, 1400, 2600$ and 4000 s/mm², with 8, 12, 28, 52 and 80 samples per shell respectively (180 gradient directions in total), in this case with AP/PA phase encoding.

3) *Preprocessing*: All data were preprocessed with a patch-based image denoising technique based on the Marchenko-Pastur distribution in the complex (phase-magnitude) image domain coupled with phase-corrected real reconstruction [30], [54], Gibbs-ringing removal [55], and motion and distortion correction [25], [56]. Dataset 1 was also corrected for signal drift due to gradient heating over the longer acquisition time [57]. Brain masking was done on the mean $b = 0$ image [58].

B. SHARD basis construction

The SHARD basis functions for multi-shell dMRI are learned from the per-shell signal coefficients in spherical harmonics. The first step in the basis construction is therefore to project each acquired shell onto the real, even-order SH basis. When the number of acquired samples varies between shells, as is the case in our data, a direct overdetermined SH fit per shell can force ℓ_{\max} to differ between shells. Although such direct fit is sufficient and entirely compatible with the theory of Section II, in this work we prefer to incorporate the Laplace-Beltrami regularizer $\frac{1}{q^2} \Delta_{S^2} \propto -\frac{1}{b} \ell(\ell+1)$ in order to enable working at fixed ℓ_{\max} across all shells, including underdetermined cases. In this way, the basis functions will span all shells for all

orders ℓ . The SH coefficients of each shell are then obtained with regularized least squares

$$\mathbf{s}_b^* = \arg \min_{\mathbf{s}} \frac{1}{n_b} \|\mathbf{y}_b - Y_b \mathbf{s}\|^2 + \gamma^2 \|\mathbf{L}_b \mathbf{s}\|^2, \quad (11)$$

where \mathbf{y}_b , \mathbf{s}_b , and n_b are respectively the signal intensities, the SH coefficients, and the number of samples in shell b . Y_b is the corresponding SH basis matrix, and \mathbf{L}_b is a diagonal matrix that contains the Laplace-Beltrami regularization factors $-\frac{1}{b} \ell(\ell+1)$. The regularizer weight γ is tuned empirically and kept small enough not to distort the power spectrum of an unregularized fit on the outer (overdetermined) shell (Fig. S.5), yielding values $\gamma = 0.02 \mu\text{m}^2/\text{ms}$ for dataset 1 and $\gamma = 0.005 \mu\text{m}^2/\text{ms}$ for dataset 2. We underscore that regularization is only used in the basis construction, not in the subsequent multi-shell fit. The SH order $\ell_{\max} = 12$ in dataset 1 and $\ell_{\max} = 8$ in dataset 2, were set to the maximum for which at least two shells are overdetermined.

The resulting SH coefficients of all shells (including $b = 0$) and all voxels within a brain mask are subsequently arranged in matrices S_ℓ with structure given by equations (3)–(4). To normalise the noise variance with respect to the number of samples in each shell, the signal is rescaled with the square root of the number of samples on each shell. To this end, a diagonal weighting matrix $\mathbf{W} = \text{diag}(\dots, \sqrt{n_b}, \dots)$ is introduced. The SVD of the reweighted matrices $\tilde{S}_\ell = \mathbf{W} S_\ell = \tilde{\mathbf{U}}_\ell \Sigma_\ell \mathbf{V}_\ell$

provides the basis matrices $U_\ell = W^{-1} \tilde{U}_\ell$ for radial q -space. These correspond to the eigenvectors of the weighted covariance matrices $W S_\ell S_\ell^\top W^\top$. As shown in Fig. 1c, this normalisation ensures uniform variance across all shells under independent simulated noise, regardless of the number of samples per shell.

We explore two possible strategies for selecting rank- r subsets of the complete basis \mathcal{U} , based on the observation that the strongest signal contribution originates from low- ℓ and low- n components. In the first strategy, the number of basis functions per harmonic order is matched to the Bessel and SHORE bases (next section) by selecting components in upper-left triangles in Fig. 2, for example using 3 basis vectors at $\ell = 0$, 2 at $\ell = 2$ and 1 at $\ell = 4$. Since each component in band ℓ contributes $2\ell + 1$ SH phase terms, in this example the total basis rank $r = 3 \times 1 + 2 \times 5 + 1 \times 9 = 22$. We will refer to this strategy as *matched* ordering. In the second strategy, the basis functions are optimally ordered for decreasing effect size ε_j as described in Sec. II and illustrated in Suppl. Tables S.1–S.2. We will refer to this strategy as *optimal* ordering.

C. Comparison with Bessel and SHORE bases

As explained in Sec. II, SHARD learns the orthonormal rank- r basis that best represents the q -space signal in the data at hand. Here, we wish to verify if SHARD indeed outperforms alternative bases and by what margin. To this end, we evaluate and compare the residual root-mean-squared error (RMSE) in a rank- r fit with equal rank representations in the spherical Bessel function basis and the SHORE basis.

The spherical Bessel functions $j_\ell(\alpha q)$ provide a basis for the Fourier transform in spherical coordinates. Here, we discretize the radial frequency α using a zero-derivative boundary condition at q_{\max} :

$$R_\ell^n(q) = \frac{1}{N_\ell^n} j_\ell(\alpha_\ell^n q) \quad , \quad (12)$$

where $\alpha_\ell^n = x_\ell^n / q_{\max}$ with x_ℓ^n the positive zeros of $j'_\ell(x)$ and N_ℓ^n is a normalization factor. Basis functions are ordered for increasing energy $\propto \alpha_\ell^n$ [48]. For the low basis rank used in our experiments, this minimum-energy ordering comes down to selecting $2(n-1) + \ell \leq \ell_{\max}$.

The SHORE basis is defined using the associated Laguerre polynomials $L_n^p(x)$:

$$R_\ell^n(q) = \frac{1}{N_\ell^n} \left(\frac{q^2}{\zeta}\right)^{\ell/2} \exp\left(\frac{-q^2}{2\zeta}\right) L_{n-\ell}^{\ell+1/2}\left(\frac{q^2}{\zeta}\right) \quad , \quad (13)$$

where N_ℓ^n is a normalization factor [40]. Here, we calibrate the scale parameter ζ to each dataset using non-linear minimization of the RMSE across the full brain mask, in order to ensure the best possible data representation. For real, symmetric signals, the index n ranges from $\ell/2$ to $\ell_{\max}/2$ [34].

D. Rank selection

SHARD defines the effect size of each basis component, but does not determine what effect size is relevant in a given application. A higher basis rank will yield a more accurate representation of the data, at the expense of having more parameters to fit. Selecting the optimal trade-off will

inevitably depend on the application. Here, we primarily consider applications that require dMRI contrast prediction, such as outlier rejection and motion correction. In this case, one possible strategy for rank selection is leave-one-out cross-validation, in which we verify the capacity of a basis of certain rank r to accurately predict unseen data samples. We then select the basis rank for which the prediction error is minimal.

E. Example application in outlier rejection

We illustrate a potential application of the SHARD basis for outlier robust regression from M-estimator theory [59], using iteratively reweighted least squares fitting with Soft- L_1 and Cauchy loss functions [60]. This scheme starts with an initial least-squares estimate in each voxel, and iteratively downweights samples with large residuals to reduce the effect of outliers. The sample weights are calculated as $w(e) = \rho'(e)/e$, according to a chosen loss function

$$\rho(e) = \begin{cases} e^2 & L_2\text{-loss} \\ 2\sqrt{1+e^2} - 2 & \text{Soft-}L_1 \text{ loss} \\ \log(1+e^2) & \text{Cauchy loss} \end{cases} \quad , \quad (14)$$

where e is the sample residual, rescaled by a normalisation factor set to the median absolute deviation in the sample.

We evaluate the performance of each loss function in low-rank multi-shell representations in data samples with varying amount of simulated outliers. A random subsample of 1000 voxels was extracted from the data in the brain mask. For any percentage of outliers, ranging from 0% to 25%, outliers are simulated by setting a randomly chosen subset of the samples in each voxel to zero. As such, we mimic at the individual voxel level the effect of slice dropouts, arguably the predominant source of outliers in EPI. We then evaluate the RMSE between the uncorrupted data samples and their predictions in linear and robust least-squares regression.

IV. RESULTS

A. SHARD basis construction and evaluation

The SHARD basis functions are plotted in Fig. 2 for dataset 1 and in Suppl. Fig. S.2 for dataset 2. These basis functions capture the covariance of the data between shells. Within each harmonic band ℓ , subsequent components capture higher radial frequencies of decreasing effect size. Across harmonic bands, the covariance between shells reduces with increasing ℓ . At $\ell \geq 10$, this results in near-independent shells: $\tilde{S}_\ell \tilde{S}_\ell^\top$ reduces to a near-diagonal matrix and basis functions for $n \geq 2$ reduce to Dirac δ -functions. The effect sizes of all components (Suppl. Tables S.1–S.2) are largest for the basis functions of low ℓ and n , plotted in the upper-left triangle in Fig. 2. A similar observation can be made in the magnitude images of all components (Suppl. Fig. S.1 and S.3).

We evaluate the RMSE of the signal fit in the matched and optimal SHARD bases, normalised to the RMS power of the measured signal, and compare this measure to other bases. Figure 3 plots the relative RMSE as a function of the basis rank r . As expected, the RMSE decreases with increasing r . We observe that SHARD outperforms the calibrated SHORE

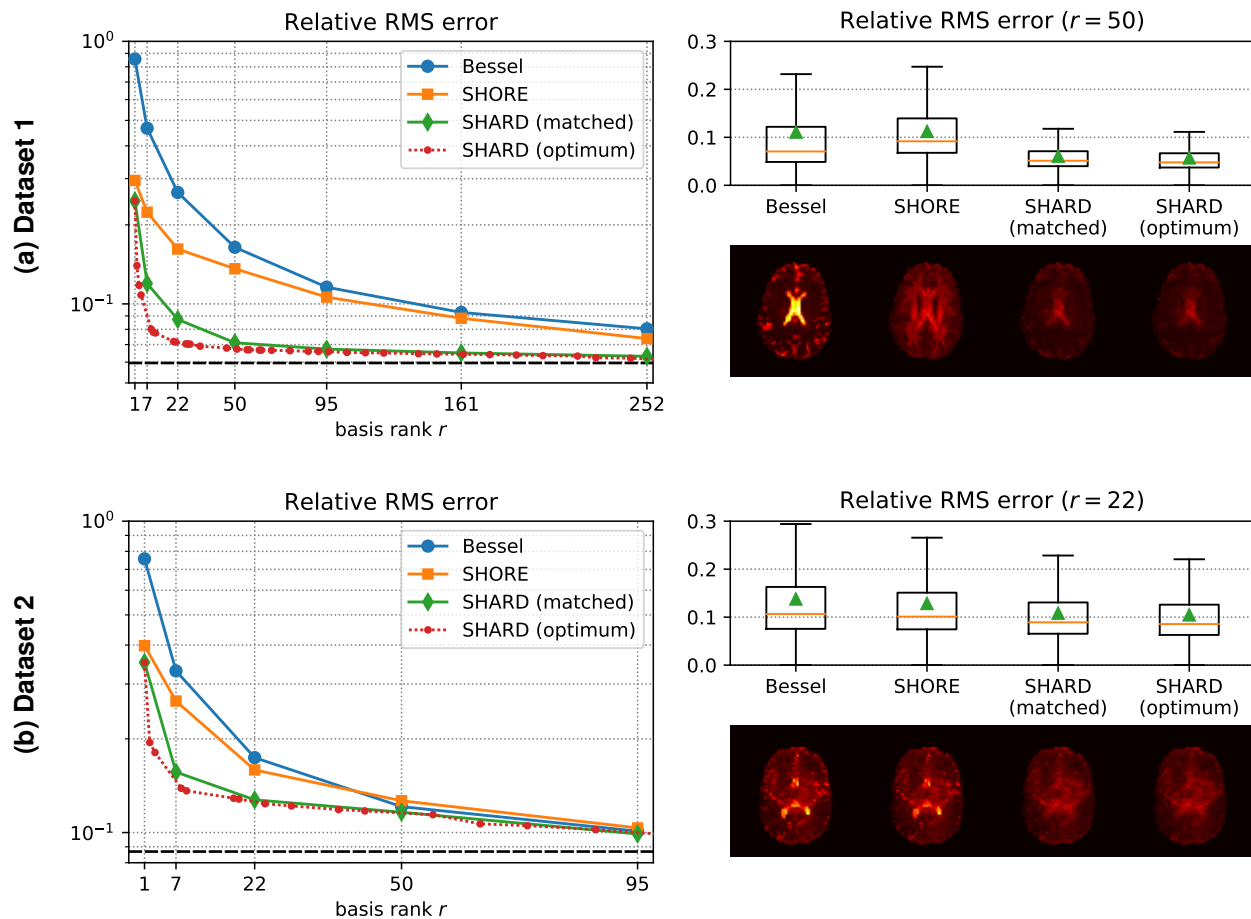


Fig. 3. Root mean squared error (RMSE) of the signal fit in the different bases, shown for dataset 1 (a) and 2 (b). The graphs on the left depict the total RMSE across the brain mask, relative to the signal power, on a logarithmic scale. This fitting error reduces for increasing basis rank r . For SHARD, the RMSE converges to the residual fitting error of an independent SH fit, indicated by the black dashed line in the graph. On the right, box plots of the voxel-wise relative RMSE are shown at rank $r = 50$ for dataset 1 and at rank $r = 22$ for dataset 2, as well as images of its value across the brain. The green triangles indicate the mean RMSE.

and Bessel bases at any rank, and most strongly in the low-rank range ($r < 50$). In dataset 1, we also found reduced residuals at high basis rank, most likely due to stronger non-monoexponential signal decay in the extended b -value range, which is more efficiently captured by SHARD. Towards high rank, the RMSE in SHARD converges to the residual fitting error of the independent SH fit from which the basis was derived. This residual SH fitting error is indicated by the black dashed line and sets a lower bound on the SHARD RMSE. We also see that the ordering strategy of the SHARD basis functions has a minor effect on accuracy in the low-rank regime, with selection according to decreasing component effect size slightly outperforming the matched order. The images on the right in Fig. 3 show that highest fitting errors occur in the ventricles due to physiological noise, and also in regions affected by an unfolding artefact in dataset 2. SHARD provides more accurate signal representation with lower and more homogenous fitting errors across the brain.

B. Rank selection

Leave-one-out cross-validation in the optimally-ordered SHARD basis selected rank $r = 46$ in dataset 1 and $r = 19$ in dataset 2 (Suppl. Tables S.1–S.2). When the component order is matched to the SHORE and Bessel bases, these values are rounded to $r = 50$ for dataset 1 and $r = 22$ for dataset 2. The best dMRI contrast prediction is thus obtained at fairly low rank, where SHARD is shown to outperform alternatives.

In addition, we compare SHARD to the per-shell SH basis by exploring the prediction error in nested leave-one-out cross-validation of single-shell and multi-shell predictions. In the SH basis, we generate a prediction for every dMRI volume from all other volumes in the same shell. In the SHARD basis, the prediction for each volume is generated from all other dMRI volumes across all shells. In both cases, the basis rank is selected to yield lowest overall prediction error per shell using cross-validation on the remaining dMRI volumes. As such, this experiment compares the best achievable prediction error in a single-shell setup, to the minimal prediction error in a multi-shell setup. The plots in Fig. 4 demonstrate that a multi-shell prediction consistently outperforms a single-shell approach,

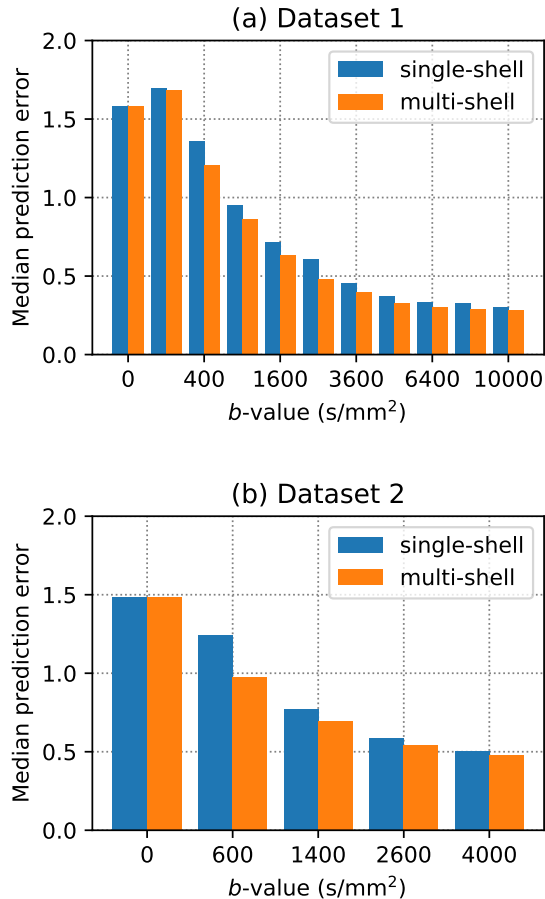


Fig. 4. Comparison of the single-shell and multi-shell prediction error in nested leave-one-out cross-validation for datasets 1 (a) and 2 (b). The blue bars show the median prediction error for single-shell SH cross-validation, where each dMRI volume is predicted from all other volumes in the shell. The orange bars show the median prediction error for multi-shell SHARD cross-validation, where each dMRI volume is predicted from all other volumes across all shells.

with improvements ranging from 4% to 20%. This shows the potential interest of using multi-shell signal representations in motion and distortion correction applications.

The low-rank projected data inevitably sacrifices part of the information in the input. We qualitatively investigate this information loss in rank-reduced data using unsupervised multi-tissue spherical factorization under convexity and nonnegativity constraints into 3 components of SH order $(8, 0, 0)$, respectively associated with white matter, grey matter, and cerebrospinal fluid [23]. The resulting fibre ODFs are depicted in Fig. 5 for the full data and for rank- r projected data. We observe that the data quality indeed improves for increasing rank. A minimum rank $r = 22$ (corresponding with SH order $\ell_{\max} = 4$) is required for a faithful data representation in crossing fibre regions. From rank $r \geq 50$, the ODF factorization becomes nearly indistinguishable from the results in the original data.

C. Outlier reweighting

The outlier robust fitting scheme using iteratively reweighted least squares with Soft- L_1 and Cauchy loss functions is

evaluated and compared to linear least squares in data with simulated outliers. The results are shown in Fig. 6 for samples drawn uniformly from dataset 1, and in Suppl. Fig. S.4 for samples drawn from dataset 2. As expected, the RMSE of the linear least squares fit increases with increasing levels of simulated outliers. In accordance with the results of Fig. 3, SHARD outperforms the Bessel and SHORE bases at low outlier levels. At outlier levels $\geq 20\%$, the margin narrows and all representations are equally corrupted. The robust estimator with Soft- L_1 loss substantially improves the RMSE at outlier levels over 5%, at the expense of a small performance penalty in uncorrupted data. This improvement is stronger in SHARD than in the alternative representations. The benefit of a Cauchy loss function is comparatively smaller in the Bessel and SHORE representations and slightly larger in SHARD, widening the performance margin between them. In all experiments, SHARD achieved higher accuracy in outlier robust fitting than the other representations.

In addition, we validated the sample weights, assigned during iteratively reweighted least squares, in relation to the simulated outlier mask using receiver operator characteristic (ROC) curves shown in the bottom row of Figs. 6 and S.4. ROC curves plot sensitivity versus specificity, parametrized by the weight threshold value. The area under the ROC curve, which provides an aggregate measure of classifier performance, is highest with the SHARD basis in all experiments, indicating improved capability to discriminate outlier samples.

V. DISCUSSION

The SHARD basis provides a linear, orthonormal representation of q -space dMRI data, with potential interest for a wide range of techniques and applications that exploit redundancy or rank-reduction. This decomposition builds on minimal prior assumptions, namely antipodal symmetry of the q -space signal and isotropic (rotation-invariant) frequency spectra of the radial basis. In contrast to the SHORE and Bessel function bases, this data representation is calibrated to the data at hand, akin to other data-driven and blind source separation techniques [23]. Our results have shown that this data-driven basis construction reduces the residuals in fitting multi-shell data, especially in the low-rank regime. SHARD hence provides a more accurate data representation for matched number of parameters.

As shown in Section II, the resulting basis spans the complete microstructural information encoded in the signal and can thus provide a lossless data representation. Its primary advantage, however, is for compact, lossy data representation. This property was verified experimentally by the two-fold RMSE reduction at low rank in Fig. 3, and further illustrated for unsupervised multi-tissue decomposition in Fig. 5, showing that the low-rank representation facilitates highly efficient data compression. SHARD hence provides a suitable and efficient characterization of the measured dMRI signal and the underlying microstructure.

A closer examination of Fig. 2 reveals first of all that the covariance matrices qualitatively resemble those of simulated white matter in Fig. 1b, though not of rank-1 and with an additional imprint of noise on the diagonal. Secondly, the derived basis functions, whilst obtained without any pre-imposed model, have desirable physical properties also found

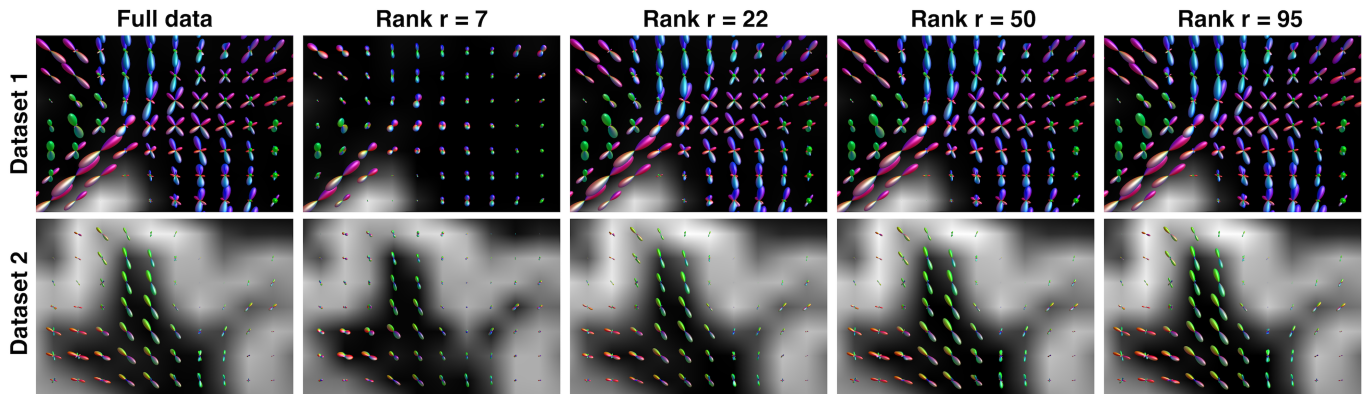


Fig. 5. Fibre orientation distribution functions in the full data (left column), and rank-reduced data in the matched SHARD basis. The top row depicts the centrum semiovale in dataset 1, overlaid onto the CSF tissue fraction. The bottom row shows the superior frontal gyrus in dataset 2, overlaid onto the grey matter fraction.

in the SHORE basis: the $n = 1$ components of $\ell = 0, 2, 4, \dots$ capture continuously higher b -values and the $n > 1$ components behave as orthogonal harmonics of increasing frequency. The SH regularization in equation (11), although not technically needed, can help reducing adverse effects from physiological noise in the $b \leq 400 \text{ s/mm}^2$ shells (Suppl. Fig. S.5) that are otherwise captured in the basis functions. Its weight γ thus offers control over unphysical (non-diffusion) effects in the data-driven basis construction that can be explored in relation to each application.

The main limitation of the current formalism is the requirement for spherically sampled data. This limitation stems from the fact that SHARD decomposes the radial q -space domain independently after projecting each shell to the SH basis. Spline interpolation or other numerical techniques can nevertheless be used to predict the signal between shells, either by interpolating the basis vectors in (9) or by interpolating the rows of matrix (3) before decomposition. Future work can investigate means of extending the signal representation to arbitrary q -space sampling. This would require extending the current radial decomposition to the joint radial and angular domain, inherently encompassing the spherical harmonics projection. In addition, a joint groupwise data representation currently requires all subjects to be acquired with the same b -value sampling scheme, due to the discrete nature of the basis functions. However, when this requirement is fulfilled, a common group-level basis could be directly derived by extending matrix (3) with voxel data of multiple subjects. This too can be subject of future work.

The SHARD representation has numerous potential applications for dMRI processing and analysis. This work has already illustrated a first, straightforward example in outlier reweighting using robust regression. Outlier robust fitting exploits optimal rank-reduction of the SVD, leveraging oversampling in q -space to detect outliers that lie too far from their predicted values. One can then either reject or downweight these outlier samples in further analysis, or proceed directly with the robustly-fitted SHARD signal representation. Another potential application is motion and distortion correction, in which data of individual volumes or slices are aligned with a common, iteratively-updated registration target, usually obtained as a low-rank

prediction of the data. A proof-of-concept of using SHARD in motion correction was recently shown in [61]. The results of the leave-one-out cross-validation in Fig. 4 indicate that it is indeed advantageous to predict individual slices and volumes from all data, rather than from a single-shell.

Furthermore, it is worth noting that the advantages of SHARD for dimensionality reduction can also be applied across the radial q -domain. This can be a useful property for selecting the number of shells and their most discriminating b -values in multi-shell protocol design, as was already successfully demonstrated on isotropic ($\ell = 0$) averages [49]. SHARD can extend this approach to incorporate the higher harmonic orders.

Finally, it has not escaped our notice that the rank-1 decomposition of Eq. (7) suggests a generalised set of rotation-invariant signal features at the voxel or patch level. Indeed, these local *SHARD features* may offer a potential means for predicting microstructure parameters, akin to recent methods that rely on signal features extracted per shell [19], [20], whilst exploiting the covariance in the combined multi-shell data and thus potentially less prone to noise. A patch-level SHARD decomposition may also be combined with random matrix theory for determining the optimal rank threshold in image denoising [30]. Future work can explore the merits of these extended applications.

VI. CONCLUSION

SHARD provides a complete, orthogonal representation of the dMRI signal, tailored to the spherical geometry of q -space and calibrated to the data at hand. Rank-reduced SHARD decomposition outperformed model-based alternatives tested, whilst maximally capturing the micro- and meso-structural information in the signal. As such, SHARD is better suited for applications that require low-rank data predictions, such as outlier rejection and motion correction.

REFERENCES

- [1] D. Le Bihan, E. Breton, D. Lallemand, P. Grenier, E. Cabanis, and M. Laval-Jeantet, "MR imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders." *Radiology*, vol. 161, no. 2, pp. 401–407, Nov 1986.

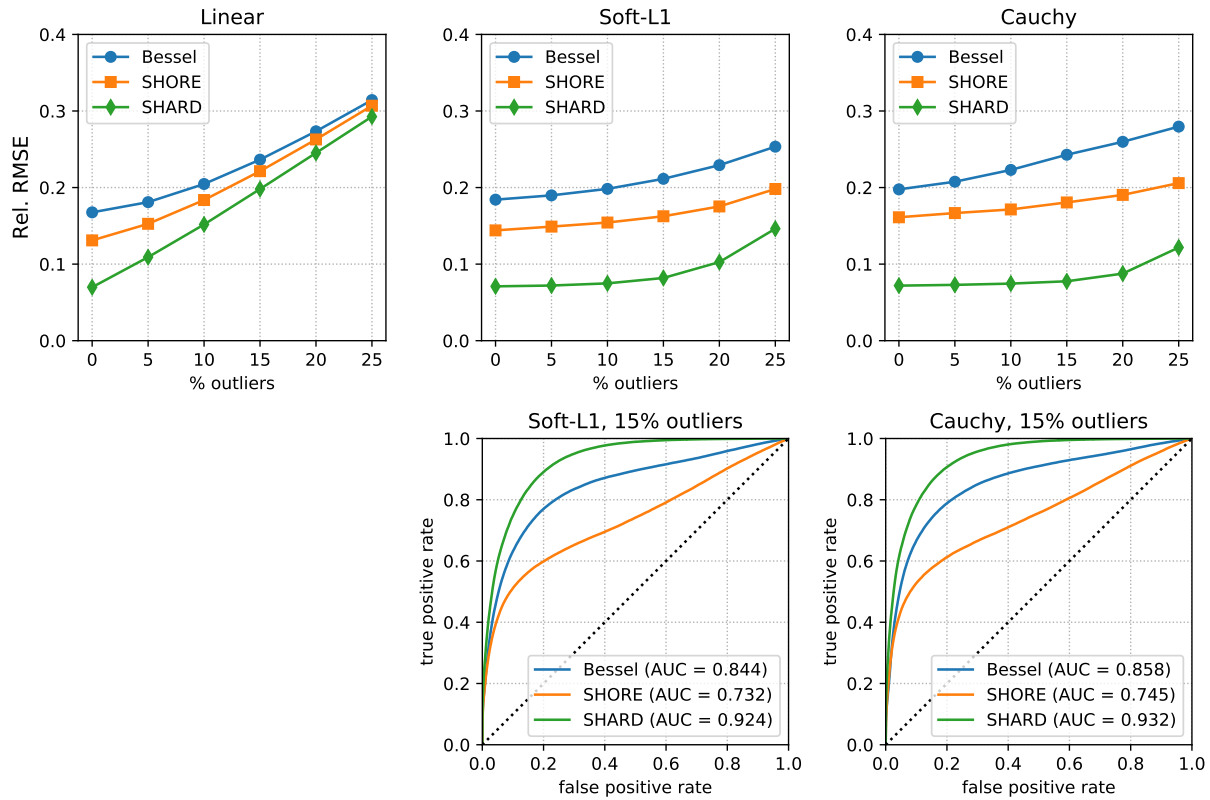


Fig. 6. Performance of robust fitting under different simulated outlier probabilities in dataset 1 using basis rank $r = 50$. Graphs in the top row show the relative RMSE between the original data and the predictions after linear least squares fitting, and after iterative reweighted least squares with Soft- L_1 and Cauchy loss functions. The bottom row plots receiver operator characteristic (ROC) curves of the sample outlier weights in the robust fitting at a simulated outlier level of 15%.

- [2] P. T. Callaghan, C. D. Eccles, and Y. Xia, "NMR microscopy of dynamic displacements: k-space and q-space imaging," *Journal of Physics E: Scientific Instruments*, vol. 21, no. 8, pp. 820–822, aug 1988.
- [3] L. R. Frank, "Characterization of anisotropy in high angular resolution diffusion-weighted MRI," *Magnetic Resonance in Medicine*, vol. 47, no. 6, pp. 1083–1099, jun 2002.
- [4] D. S. Tuch, T. G. Reese, M. R. Wiegell, N. Makris, J. W. Belliveau, and V. J. Wedeen, "High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity," *Magnetic Resonance in Medicine*, vol. 48, no. 4, pp. 577–582, sep 2002.
- [5] D. S. Tuch, "Q-ball imaging," *Magnetic Resonance in Medicine*, vol. 52, no. 6, pp. 1358–1372, 2004.
- [6] A. W. Anderson, "Measurement of fiber orientation distributions using high angular resolution diffusion imaging," *Magnetic Resonance in Medicine*, vol. 54, no. 5, pp. 1194–1206, 2005.
- [7] J.-D. Tournier, F. Calamante, D. G. Gadian, and A. Connelly, "Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution," *NeuroImage*, vol. 23, no. 3, pp. 1176–1185, nov 2004.
- [8] J.-D. Tournier, F. Calamante, and A. Connelly, "Robust determination of the fibre orientation distribution in diffusion MRI: Non-negativity constrained super-resolved spherical deconvolution," *NeuroImage*, vol. 35, no. 4, pp. 1459–1472, May 2007.
- [9] M. Descoteaux, R. Deriche, T. Knosche, and A. Anwander, "Deterministic and probabilistic tractography based on complex fibre orientation distributions," *IEEE Transactions on Medical Imaging*, vol. 28, no. 2, pp. 269–286, Feb 2009.
- [10] Y. Assaf and P. J. Basser, "Composite hindered and restricted model of diffusion (CHARMED) MR imaging of the human brain," *NeuroImage*, vol. 27, no. 1, pp. 48–58, aug 2005.
- [11] Y. Assaf, T. Blumenfeld-Katzir, Y. Yovel, and P. J. Basser, "Axcaliber: A method for measuring axon diameter distribution from diffusion MRI," *Magnetic Resonance in Medicine*, vol. 59, no. 6, pp. 1347–1354, 2008.
- [12] D. C. Alexander, P. L. Hubbard, M. G. Hall, E. A. Moore, M. Ptito, G. J. Parker, and T. B. Dyrby, "Orientationally invariant indices of axon diameter and density from diffusion MRI," *NeuroImage*, vol. 52, no. 4, pp. 1374–1389, oct 2010.
- [13] E. Fieremans, J. H. Jensen, and J. A. Helpert, "White matter characterization with diffusional kurtosis imaging," *NeuroImage*, vol. 58, no. 1, pp. 177–188, sep 2011.
- [14] H. Zhang, T. Schneider, C. A. Wheeler-Kingshott, and D. C. Alexander, "NODDI: Practical in vivo neurite orientation dispersion and density imaging of the human brain," *NeuroImage*, vol. 61, no. 4, pp. 1000–1016, jul 2012.
- [15] S. N. Jespersen, L. A. Leigland, A. Cornea, and C. D. Kroenke, "Determination of axonal and dendritic orientation distributions within the developing cerebral cortex by diffusion tensor imaging," *IEEE Transactions on Medical Imaging*, vol. 31, no. 1, pp. 16–32, jan 2012.
- [16] M. Reisert, V. G. Kiselev, B. Dital, E. Kellner, and D. S. Novikov, "MesoFT: Unifying diffusion modelling and fiber tracking," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2014*, ser. Lecture Notes in Computer Science, P. Golland, N. Hata, C. Barrilot, J. Hornegger, and R. Howe, Eds. Springer International Publishing, 2014, vol. 8675, pp. 201–208.
- [17] A. Daducci, E. J. Canales-Rodríguez, H. Zhang, T. B. Dyrby, D. C. Alexander, and J.-P. Thiran, "Accelerated microstructure imaging via convex optimization (AMICO) from diffusion MRI data," *NeuroImage*, vol. 105, pp. 32–44, jan 2015.
- [18] E. Kaden, F. Kruggel, and D. C. Alexander, "Quantitative mapping of the per-axon diffusion coefficients in brain white matter," *Magnetic Resonance in Medicine*, vol. 75, no. 4, pp. 1752–1763, apr 2016.
- [19] M. Reisert, E. Kellner, B. Dital, J. Hennig, and V. G. Kiselev, "Disentangling micro from mesostructure by diffusion MRI: A bayesian approach," *NeuroImage*, vol. 147, pp. 964–975, feb 2017.
- [20] D. S. Novikov, J. Veraart, I. O. Jelescu, and E. Fieremans, "Rotationally-invariant mapping of scalar and orientational metrics of neuronal microstructure with diffusion MRI," *NeuroImage*, vol. 174, pp. 518–538, jul 2018.

- [21] B. Jeurissen, J.-D. Tournier, T. Dhollander, A. Connelly, and J. Sijbers, "Multi-tissue constrained spherical deconvolution for improved analysis of multi-shell diffusion MRI data," *NeuroImage*, vol. 103, pp. 411–426, Dec 2014.
- [22] D. Christiaens, M. Reisert, T. Dhollander, S. Sunaert, P. Suetens, and F. Maes, "Global tractography of multi-shell diffusion-weighted imaging data using a multi-tissue model," *NeuroImage*, vol. 123, pp. 89–101, dec 2015.
- [23] D. Christiaens, S. Sunaert, P. Suetens, and F. Maes, "Convexity-constrained and nonnegativity-constrained spherical factorization in diffusion-weighted imaging," *NeuroImage*, vol. 146, pp. 507–517, feb 2017.
- [24] G. Rohde, A. Barnett, P. Basser, S. Marengo, and C. Pierpaoli, "Comprehensive approach for correction of motion and distortion in diffusion-weighted MRI," *Magnetic Resonance in Medicine*, vol. 51, no. 1, pp. 103–114, 2004.
- [25] J. L. Andersson and S. N. Sotiropoulos, "An integrated approach to correction for off-resonance effects and subject movement in diffusion MR imaging," *NeuroImage*, vol. 125, pp. 1063–1078, jan 2016.
- [26] L.-C. Chang, D. K. Jones, and C. Pierpaoli, "RESTORE: Robust estimation of tensors by outlier rejection," *Magnetic Resonance in Medicine*, vol. 53, no. 5, pp. 1088–1095, 2005.
- [27] K. Pannek, D. Raffelt, C. Bell, J. L. Mathias, and S. E. Rose, "HOMOR: Higher order model outlier rejection for high b-value MR diffusion data," *NeuroImage*, vol. 63, no. 2, pp. 835–842, nov 2012.
- [28] A. Tobisch, T. Stocker, S. Groeschel, and T. Schultz, "Iteratively reweighted L1-fitting for model-independent outlier removal and regularization in diffusion MRI," in *2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI)*. IEEE, apr 2016.
- [29] B. Scherrer, A. Gholipour, and S. K. Warfield, "Super-resolution reconstruction to increase the spatial resolution of diffusion weighted images from orthogonal anisotropic acquisitions," *Medical Image Analysis*, vol. 16, no. 7, pp. 1465–1476, oct 2012.
- [30] J. Veraart, D. S. Novikov, D. Christiaens, B. Ades-aron, J. Sijbers, and E. Fieremans, "Denoising of diffusion MRI using random matrix theory," *NeuroImage*, vol. 142, pp. 394–406, nov 2016.
- [31] H. Mirzaalian, L. Ning, P. Savadjiev, O. Pasternak, S. Bouix, O. Michailovich, G. Grant, C. Marx, R. Morey, L. Flashman, M. George, T. McAllister, N. Andaluz, L. Shutter, R. Coimbra, R. Zafonte, M. Coleman, M. Kubicki, C. Westin, M. Stein, M. Shenton, and Y. Rathi, "Inter-site and inter-scanner diffusion MRI data harmonization," *NeuroImage*, vol. 135, pp. 311–323, jul 2016.
- [32] M. Reisert, H. Skibbe, and V. G. Kiselev, "The diffusion dictionary in the human brain is short: Rotation invariant learning of basis functions," in *Computational Diffusion MRI and Brain Connectivity*, ser. Mathematics and Visualization, T. Schultz, G. Nedjati-Gilani, A. Venkataraman, L. O'Donnell, and E. Panagiotaki, Eds. Springer, New York, 2014, pp. 47–55.
- [33] D. Christiaens, F. Maes, S. Sunaert, and P. Suetens, "Convex non-negative spherical factorization of multi-shell diffusion-weighted images," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, ser. Lecture Notes in Computer Science, N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, Eds. Springer International Publishing, 2015, vol. 9349, pp. 166–173.
- [34] E. Özarslan, C. G. Koay, T. M. Shepherd, M. E. Komlosh, M. O. İrfanoğlu, C. Pierpaoli, and P. J. Basser, "Mean apparent propagator (MAP) MRI: A novel diffusion imaging method for mapping tissue microstructure," *NeuroImage*, vol. 78, pp. 16–32, sep 2013.
- [35] R. H. Fick, D. Wassermann, E. Caruyer, and R. Deriche, "MAPL: Tissue microstructure estimation using laplacian-regularized MAP-MRI and its application to HCP data," *NeuroImage*, vol. 134, pp. 365–385, jul 2016.
- [36] H.-E. Assemlal, D. Tschumperlé, and L. Brun, "Efficient and robust computation of PDF features from diffusion MR signal," *Medical Image Analysis*, vol. 13, no. 5, pp. 715–729, oct 2009.
- [37] M. Descoteaux, R. Deriche, D. L. Bihan, J.-F. Mangin, and C. Poupon, "Multiple q-shell diffusion propagator imaging," *Medical Image Analysis*, vol. 15, no. 4, pp. 603–621, aug 2011.
- [38] E. Caruyer and R. Deriche, "Diffusion MRI signal reconstruction with continuity constraint and optimal regularization," *Medical Image Analysis*, vol. 16, no. 6, pp. 1113–1120, aug 2012.
- [39] S. Merlet, J. Cheng, A. Ghosh, and R. Deriche, "Spherical polar fourier EAP and ODF reconstruction via compressed sensing in diffusion MRI," in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. IEEE, mar 2011, pp. 365–371.
- [40] S. L. Merlet and R. Deriche, "Continuous diffusion signal, EAP and ODF estimation via compressive sensing in diffusion MRI," *Medical Image Analysis*, vol. 17, no. 5, pp. 556–572, jul 2013.
- [41] A. P. Hosseinbor, M. K. Chung, Y.-C. Wu, and A. L. Alexander, "Bessel fourier orientation reconstruction (BFOR): An analytical diffusion propagator reconstruction for hybrid diffusion imaging and computation of q-space indices," *NeuroImage*, vol. 64, pp. 650–670, jan 2013.
- [42] Y. Rathi, O. Michailovich, F. Laun, K. Setsompop, P. Grant, and C.-F. Westin, "Multi-shell diffusion signal recovery from sparse measurements," *Medical Image Analysis*, vol. 18, no. 7, pp. 1143–1156, oct 2014.
- [43] V. G. Kiselev, "The cumulant expansion: an overarching mathematical framework for understanding diffusion NMR," in *Diffusion MRI*, D. K. Jones, Ed. Oxford University Press, 2010, pp. 152–168.
- [44] E. Özarslan and T. H. Mareci, "Generalized diffusion tensor imaging and analytical relationships between diffusion tensor imaging and high angular resolution diffusion imaging," *Magnetic Resonance in Medicine*, vol. 50, no. 5, pp. 955–965, oct 2003.
- [45] C. Liu, R. Bammer, B. Acar, and M. E. Moseley, "Characterizing non-gaussian diffusion by using generalized diffusion tensors," *Magnetic Resonance in Medicine*, vol. 51, no. 5, pp. 924–937, 2004.
- [46] T. Schultz, A. Fuster, A. Ghosh, R. Deriche, L. Florack, and L.-H. Lim, "Higher-order tensors in diffusion imaging," in *Visualization and Processing of Tensors and Higher Order Descriptors for Multi-Valued Data*, ser. Mathematics and Visualization, C.-F. Westin, A. Vilanova, and B. Burgeth, Eds. Springer, 2014, pp. 129–161.
- [47] J. L. Andersson and S. N. Sotiropoulos, "Non-parametric representation and prediction of single- and multi-shell diffusion-weighted MRI data using gaussian processes," *NeuroImage*, vol. 122, pp. 166–176, nov 2015.
- [48] Q. Wang, O. Ronneberger, and H. Burkhardt, "Rotational invariance based on fourier analysis in polar and spherical coordinates," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 9, pp. 1715–1722, sep 2009.
- [49] J.-D. Tournier, E. Hughes, N. Tumor, S. N. Sotiropoulos, S. Jbabdi, J. Andersson, D. Rueckert, A. D. Edwards, and J. V. Hajnal, "Data-driven optimisation of multi-shell HARDI," in *ISMRM Annual Meeting & Exhibition*, vol. 23, may 2015, p. 2897.
- [50] S. N. Jespersen, C. D. Kroenke, L. Østergaard, J. J. Ackerman, and D. A. Yablonskiy, "Modeling dendrite density from magnetic resonance diffusion measurements," *NeuroImage*, vol. 34, no. 4, pp. 1473–1486, feb 2007.
- [51] M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz, "Rotation invariant spherical harmonic representation of 3D shape descriptors," in *Eurographics Symposium on Geometry Processing*, L. Kobbelt, P. Schröder, and H. Hoppe, Eds., jun 2013, pp. 1–9.
- [52] J. Hutter, J. D. Tournier, A. N. Price, L. Cordero-Grande, E. J. Hughes, S. Malik, J. Steinweg, M. Bastiani, S. N. Sotiropoulos, S. Jbabdi, J. Andersson, A. D. Edwards, and J. V. Hajnal, "Time-efficient and flexible design of optimized multishell HARDI diffusion," *Magnetic Resonance in Medicine*, vol. 79, no. 3, pp. 1276–1292, may 2017.
- [53] L. Cordero-Grande, A. Price, J. Hutter, E. Hughes, and J. V. Hajnal, "Comprehensive CG-SENSE reconstruction of SMS-EPI," in *ISMRM Annual Meeting & Exhibition*, vol. 24, may 2016, p. 3239.
- [54] M. Gavish and D. L. Donoho, "Optimal shrinkage of singular values," *IEEE Transactions on Information Theory*, vol. 63, no. 4, pp. 2137–2152, apr 2017.
- [55] E. Kellner, B. Dhital, V. G. Kiselev, and M. Reisert, "Gibbs-ringing artifact removal based on local subvoxel-shifts," *Magnetic Resonance in Medicine*, vol. 76, no. 5, pp. 1574–1581, nov 2016.
- [56] J. L. Andersson, S. Skare, and J. Ashburner, "How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging," *NeuroImage*, vol. 20, no. 2, pp. 870–888, oct 2003.
- [57] S. B. Vos, C. M. W. Tax, P. R. Luijten, S. Ourselin, A. Leemans, and M. Froeling, "The importance of correcting for signal drift in diffusion MRI," *Magnetic Resonance in Medicine*, vol. 77, no. 1, pp. 285–299, jan 2016.
- [58] S. M. Smith, "Fast robust automated brain extraction," *Human Brain Mapping*, vol. 17, no. 3, pp. 143–155, nov 2002.
- [59] P. J. Huber, *Robust Statistics*. John Wiley & Sons, Inc., feb 1981.
- [60] P. W. Holland and R. E. Welsch, "Robust regression using iteratively reweighted least-squares," *Communications in Statistics - Theory and Methods*, vol. 6, no. 9, pp. 813–827, jan 1977.
- [61] D. Christiaens, L. Cordero-Grande, M. Pietsch, J. Hutter, A. D. Edwards, M. Deprez, J. V. Hajnal, and J.-D. Tournier, "Multi-shell SHARD reconstruction from scattered slice diffusion MRI data in the neonatal brain," in *Joint Annual Meeting ISMRM-ESMRMB*, vol. 26, june 2018, p. 464.

TABLE S.1

EFFECT SIZES ε_j OF ALL COMPONENTS $R_\ell^n(q)$ IN SORTED ORDER, AND THE CUMULATIVE BASIS RANK FOR DATASET 1. THE ASTERISK IN THE LAST COLUMN INDICATES THE OPTIMAL RANK AS SELECTED WITH LEAVE-ONE-OUT CROSS-VALIDATION.

ℓ	n	effect $\sqrt{\varepsilon_j}$	rank r
0	0	229.308	1
0	1	48.075	2
0	2	17.846	3
0	3	11.081	4
2	0	7.336	9
0	4	4.303	10
0	5	2.212	11
4	0	2.089	20
0	6	1.597	21
2	1	1.042	26
0	7	0.772	27
0	8	0.711	28
2	2	0.695	33
6	0	0.680	46 *
2	3	0.629	51
2	4	0.593	56
0	9	0.502	57
2	5	0.496	62
0	10	0.439	63
2	6	0.403	68
4	1	0.382	77
4	2	0.371	86
2	7	0.354	91
2	8	0.344	96
4	3	0.342	105
4	4	0.333	114
4	5	0.317	123
8	0	0.308	140
4	6	0.306	149
6	1	0.283	162
6	2	0.273	175
6	3	0.264	188
6	4	0.243	201
8	1	0.237	218
4	7	0.229	227
8	2	0.209	244
6	5	0.203	257

TABLE S.2

EFFECT SIZES ε_j OF ALL COMPONENTS $R_\ell^n(q)$ IN SORTED ORDER, AND THE CUMULATIVE BASIS RANK FOR DATASET 2. THE ASTERISK IN THE LAST COLUMN INDICATES THE OPTIMAL RANK AS SELECTED WITH LEAVE-ONE-OUT CROSS-VALIDATION.

ℓ	n	effect $\sqrt{\varepsilon_j}$	rank r
0	0	86.981	1
0	1	27.253	2
0	2	6.697	3
2	0	4.783	8
0	3	2.482	9
4	0	1.333	18
0	4	1.320	19 *
2	1	1.297	24
2	2	0.922	29
4	1	0.807	38
2	3	0.688	43
6	0	0.659	56
4	2	0.659	65
4	3	0.634	74
6	1	0.606	87
8	0	0.564	104




Fig. S.1. SHARD components of the multi-shell data for each harmonic band $\ell = 0, 2, \dots, 12$ in dataset 1. The spatial maps depict the RMS power across the SH phase m in each voxel. Only the 6 leading components are shown.

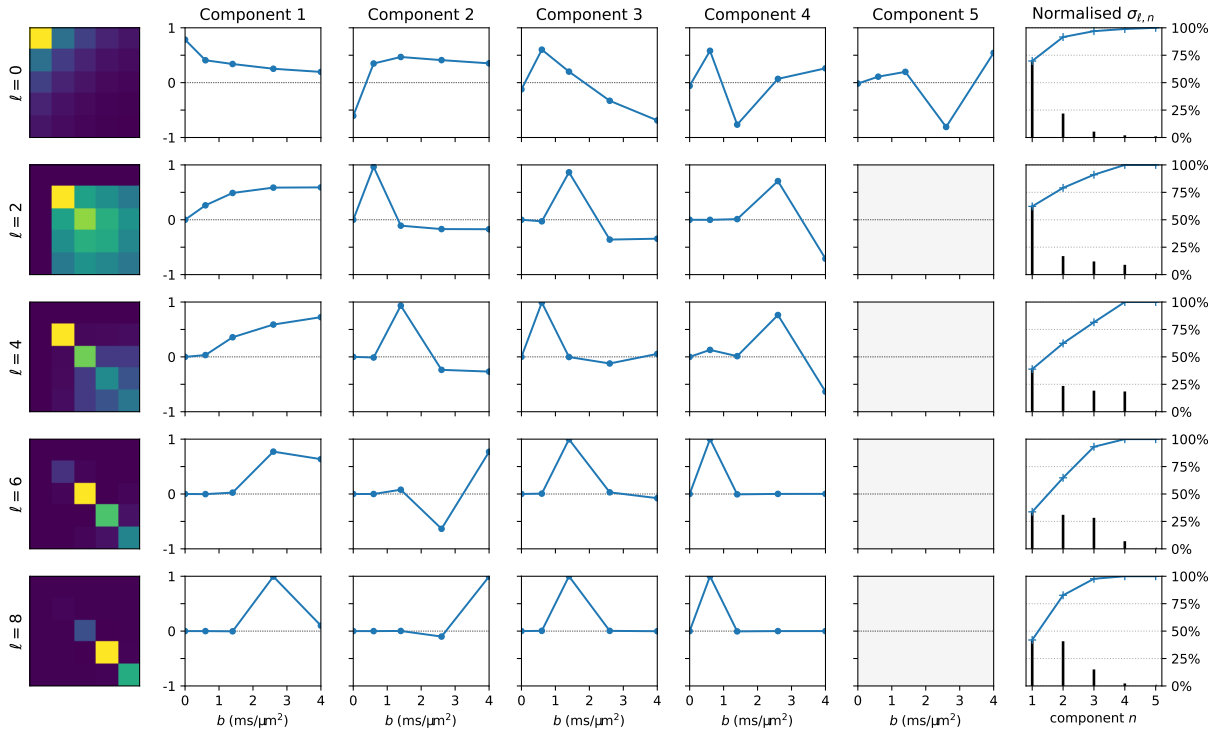


Fig. S.2. SHARD decomposition of dataset 2. Rows correspond to spherical harmonic bands of order $\ell = 0, 2, \dots, 8$. The left column illustrates square, symmetric matrices $\tilde{S}_\ell \tilde{S}_\ell^T$, showing the covariance between shells in each harmonic band. The middle columns plot the decomposition basis vectors, corresponding to the eigenvectors of the matrices on the left. Shaded boxes indicate components with $\sigma_{\ell,n} = 0$, defined only by their orthogonality to the leading (data-driven) components and excluded from basis construction. The plots in the right column depict the singular values of each component in decreasing order and their cumulative sum.




Fig. S.3. SHARD components of the multi-shell data for each harmonic band $\ell = 0, 2, \dots, 8$ in dataset 2. The spatial maps depict the RMS power across the SH phase m in each voxel.

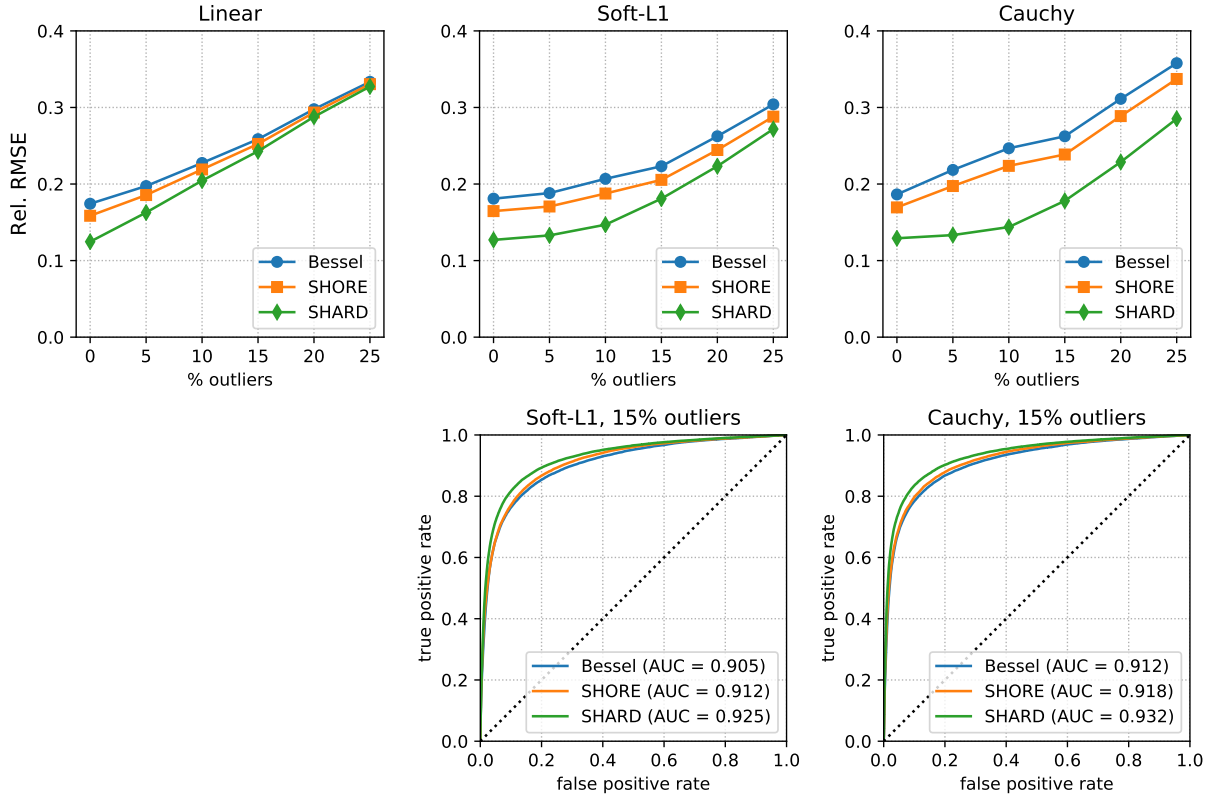


Fig. S.4. Performance of robust fitting under different simulated outlier probabilities in dataset 2 using basis rank $r = 22$. Graphs in the top row show the relative RMSE between the original data and the predictions after linear least squares fitting, and after iterative reweighted least squares with Soft- L_1 and Cauchy loss functions. The bottom row plots receiver operator characteristic (ROC) curves of the sample outlier weights in the robust fitting at a simulated outlier level of 15%.




Fig. S.5. Power spectra in multi-shell spherical harmonics (SH) fitting without and with regularization. Conventional unregularized fitting (left column) requires truncating ℓ_{\max} in each shell to the maximum overdetermined order, which can lead to the staircasing effect seen in dataset 1. Adding Laplace-Beltrami regularization (right column) enables fitting all shells at the same ℓ_{\max} . We can also observe that the regularizer effectively suppresses high-frequency components in the $b = 100 \text{ s/mm}^2$ and $b = 400 \text{ s/mm}^2$ shells of dataset 1, attributed to physiological noise rather than dMRI signal.