DOI: 10.1111/cgf.15082
Eurographics Conference on Visualization (EuroVis) 2024
W. Aigner, D. Archambault, and R. Bujack
(Guest Editors)

COMPUTER GRAPHICS *forum*
Volume 43 (2024), Number 3

Sparse q-ball imaging towards efficient visual exploration of HARDI data

D.Lei[1,2] E.Miandji[2] J.Unger[2] and I. Hotz [1]

[1]Scientific Visualization Group, Linköping University, Sweden
[2] Computer Graphics and Image Processing Group, Linköping University, Sweden

Abstract

Diffusion-weighted magnetic resonance imaging (D-MRI) is a technique to measure the diffusion of water, in biological tissues. It is used to detect microscopic patterns, such as neural fibers in the living human brain, with many medical and neuroscience applications e.g. for fiber tracking. In this paper, we consider High-Angular Resolution Diffusion Imaging (HARDI) which provides one of the richest representations of water diffusion. It records the movement of water molecules by measuring diffusion under 64 or more directions. A key challenge is that it generates high-dimensional, large, and complex datasets. In our work, we develop a novel representation that exploits the inherent sparsity of the HARDI signal by approximating it as a linear sum of basic atoms in an overcomplete data-driven dictionary using only a sparse set of coefficients. We show that this approach can be efficiently integrated into the standard q-ball imaging pipeline to compute the diffusion orientation distribution function (ODF). Sparse representations have the potential to reduce the size of the data while also giving some insight into the data. To explore the results, we provide a visualization of the atoms of the dictionary and their frequency in the data to highlight the basic characteristics of the data. We present our proposed pipeline and demonstrate its performance on 5 HARDI datasets.

1. Introduction

Diffusion-weighted magnetic resonance imaging (D-MRI) describes a class of imaging techniques measuring diffusion-related signals in biological tissues [CEX88]. The resulting data contains much valuable information about tissue microstructures, e.g., fiber structures in the brain white matter. However, the data is complex, and this information is not directly accessible and requires a careful reconstruction and representation of the diffusion signal [Lee19, CT20]. In this process, visualization has also played an important role [SV19]. Optimizing the respective imaging protocols and models to derive directional information from the raw D-MRI data is still ongoing research [CT20]. Probably the best-studied model is diffusion tensor imaging [BML94], assuming a Gaussian diffusion. While it has been used successfully in many applications and many visualizations, it also has its limitations, e.g., in areas of fiber crossings [JPGJ12]. *High angular resolution diffusion imaging* (HARDI) is one of the methods that has been introduced to overcome these challenges. Thereby, a high number of diffusion directions are measured. Thus, potentially more complex tissue structures can be resolved [MHGP*21]; however, it also increases the data size and complexity significantly. A question that has been debated since the introduction of HARDI is the required or optimal number of diffusion directions. In this context, ideas from the field of compressed sampling [RMS*11] were used to address this question. However, although the underlying assumption in compressed sampling is that the signal under consideration is sparse, the concept of sparsity has so far neither been investigated in detail nor exploited in the processing, analysis, or visualization of HARDI data.

Sparse representations provide a concise yet precise data representation, with the underlying idea of finding a coordinate transformation of the data into a sparse domain. This means finding a basis (also called a dictionary) such that the number of non-zero coefficients becomes as small as possible. It has been shown that sparse representations preserve the main features of the data well and are robust with respect to noise and small fluctuations in the signal, while reducing the required storage costs significantly [ZXY*15]. There are many examples of successful applications in signal processing and computer vision [HWG*21]. A property that makes sparse representations especially interesting in the context of medical imaging and analysis is their potential to detect semantic information [WMM*09]. As the first step towards sparse modeling of HARDI data, we analyze the main characteristics of sparse representations and investigate questions related to the sparseness of the data, the accuracy of the sparse representation, and its potential for data compression. As the central goal is to develop applications for visualization and detailed analysis, we also investigate the possibility of efficiently integrating such representations into standard analysis pipelines such as *q-ball imaging* [Tuc04] to compute the diffusion orientation distribution function (ODF). Specifically, this

means estimating the sparse representation of the ODF coefficients directly from the raw HARDI data using pre-trained dictionaries. Since in practice the reconstruction of ODFs from HARDI data often involves very large groups of patients and thus large data sizes, a key benefit of our approach is that it can significantly accelerate this process. Therefore, we present a sparsity-based pipeline that can represent and process both raw HARDI data and ODF, see Figure 2, which only requires Funk-Radon transformation at the atom-level instead of the voxel-level. Our main contributions can be summarized as follows:

- An analysis of the sparsity of HARDI data concerning accuracy and compression capabilities
- A novel sparse pipeline that can represent both raw HARDI data and ODF
- Efficient computation of ODF coefficients directly from a pre-trained dictionary on raw data

The paper is organized as follows: In Section 2, we give an overview of the related work regarding diffusion imaging, HARDI, and sparse representation for high-dimensional data. Section 3 introduces the underlying concepts used in the paper. In Section 4, we present the proposed pipeline and methods in detail. Results of experiments are presented in Section 5. Finally, in Section 6, we conclude this manuscript and address challenges in future work.

2. Related Work

This paper builds on work from two research fields, diffusion MRI (D-MRI), with more than 20 years of history, and sparse representations. This section is structured accordingly.

A: Diffusion MRI (D-MRI) has developed to become a widely used tool in neuroscience applications. It enables access to structural data in fibrous tissue, such as the white matter in the brain, which was previously unavailable [AJBTdS19], and has become a main method for noninvasive mapping of connectomes [SZ19]. In clinical research, it supports the diagnosis, e.g., of Parkinson's disease [HKDMS17] or detection of traumatic brain injuries [AAB*15]. In the area of visualization, it has significantly influenced tensor field visualization, but the contributions go beyond that. Examples include glyphs for encoding local diffusion characteristics, fiber tractography, and exploratory environments. Current activities focus on integrating uncertainty [GvdVS22] into the methods and visualizations as well as cohort comparisons [ZCH*17]. For an overview of recent developments, we refer to the article by Schultz and Vilanova [SV19]. Most of these methods are designed for Diffusion Tensor Imaging (DTI) and use a second-order tensor to approximate the diffusion orientation distribution functions (ODFs) [BML94].

There is also a large body of work that deals with the visualization of HARDI (High-Angular Resolution Diffusion Imaging) data. HARDI data provides more details about fiber distribution but is also more complex, and despite the increasing amount of information, reliable and accurate modeling of fiber orientation is still a challenging and active area of research. Modeling the signal as a sum of individual fiber compartments is one approach to estimating the number and orientation of fibers, e.g., by using constrained spherical deconvolution [SWK10]. Other approaches directly estimate the diffusion orientation probability function (ODF), e.g., with q-ball imaging [Tuc04], which is often the first step in the visualization pipeline and at the center of our work. An example is the work of Peters et al. [PPvA*09] on interactive glyph rendering for exploring HARDI data. Recent work on PDE-based compression of diffusion MR images combining image space and Q-space, presented by Jumakulyyev and Schultz [JS23], partly pursues similar goals to our work. However, their focus is on lossless compression, while we provide an approximation of the data while preserving its key characteristics. Machine learning-based methods for reconstructing fiber structures are also gaining increasing attention [KVJ*21]. These are just examples of current research trends, as a complete overview is beyond this section, and we refer to some review articles [Des15, NDH*15].

B: Sparse representations have been utilized for several tasks within the fields of image processing and computational imaging; examples include medical imaging [RB11], classification and clustering [JLD13], and image denoising [YSL13]. A related field to sparse representation is compressed sensing, where the goal is to derive provably optimal sampling algorithms for sparse signals. Compressed sensing, has been used for designing a digital camera with a single pixel [DDT*08], light field imaging [AN10], and Magnetic Resonance Imaging (MRI) [LDP07]. Schwab et al. [SVC18] have utilized compressed sensing using separable dictionaries to reduce the number of measurements for HARDI. We demarcate our method from that of Schwab et al. by noting that our focus is on the problem of sparse representation with learned dictionaries to investigate their application for identifying structures in the HARDI data, while Schwab et al. present a compressed sensing algorithm using analytical dictionaries [VRG22].

C: Compression of volume data within the field of visualization aims to enable a faster and more scalable visualization for rendering, data transmission, and storage. Balsa Rodríguez et al. [BRGIG*14] provide a comprehensive overview of existing methods in compressed GPU-based direct volume rendering (DVR). Schneider et al. [SW03] proposed hierarchical vector quantizations to represent volumetric data, Gobbetti et al. [GIGM12] utilized K-SVD to learn a dictionary for representing sampled subsets of large volumes, and Díaz et al. [DMG20] utilized quantized variable-length sparse representation for each octree-based frame. Due to the unique properties of HARDI data, e.g. the high dimensionality, these techniques cannot be directly applied without significant modification. We demarcate our method from the techniques above by noting that we propose a novel q-ball imaging pipeline rather than a HARDI data compression algorithm.

3. Background

This section provides the background for our work, the HARDI data, and the analytical q-ball imaging pipeline.

3.1. HARDI Data description

Diffusion MRI (D-MRI) describes imaging methods for measuring water diffusion in biological tissues, which enable conclusions to be drawn about the local tissue microstructure. High Angular Resolution Diffusion Imaging (HARDI) is a special variant (protocol) of D-MRI that produces diffusion-weighted volumes that represent a large number of diffusion directions. The different directions are

(a) *Example 1* **(b)** *Example 2*

Figure 1: *Two examples for the Funk-Radon Transform: The left side of each example shows the raw HARDI signal with spherical harmonics representation, and the right side shows their corresponding ODF. The color represents the corresponding value projected on the set of modified spherical harmonic bases, where red indicates a lower value and blue suggests a higher value.*

obtained by sampling the space of possible diffusion-sensitizing magnetic field gradients, spanning the so-called q-space [ST65]. In single-shell HARDI the magnitude of b is fixed and only the direction is varied [MT14]. Important parameters are the number of sampled directions and the choice of the b-value. The HARDI datasets we used were downloaded through the Laboratory of Neuro Imaging Image Data Archive (https://ida.loni.usc.edu) [F*16], collected by the Human Connectome Project (HCP) database [CCFa]. All participating subjects provided informed written consent prior to participation in the study, approved by the institutional review board approval and procedures [CCFb]. The diffusion MRI brain images were acquired using a 3 Tesla Siemens Connectome scanner. The imaging parameters for each subject were set as follows: image matrix$= 140\times 140\times 96$ and voxel resolution $= 1.5\times 1.5\times 1.5\ mm^3$, 64 and 128 diffusion-weighted volumes were acquired for each of the $b = 1000, 3000\ s/mm^2$ values and one non-diffusion weighted image was acquired ($b = 0\ s/mm^2$).

In this work, we use the 64 diffusion-weighted volumes with $b = 1000\ s/mm^2$ The data is represented as 4-D data of size $w\times h\times s\times d$, and the first three dimensions represent the spatial extent: width w, height h and a section s of the volume; The last dimension contains a vector with the measured diffusion values.

3.2. Q-ball imaging pipeline

Q-Ball Imaging (QBI) is a method for reconstructing the diffusion orientation distribution function (ODF) from the HARDI data. It has been shown that the ODF can be approximated by the Funk-Radon transform (FRT) of the diffusion signal [Tuc04]. We use a fast and robust method for estimating the ODF, which has been proposed by Descoteaux et al. [DAFD07]. The main idea is to first approximate the raw HARDI data with a set of modified spherical harmonics (SH) and then combine the FRT with the Funk-Hecke theorem, which reduces the calculation of the Funk-Radon integral over the sphere to a scaling of the SH coefficients.

3.2.1. Modified spherical harmonics for HARDI

The spherical harmonics (SH) basis function of degree l and order m, denoted Y_l^m, is defined by:

$$Y_l^m(\theta,\phi) = \sqrt{\frac{2l+1}{4\pi}\frac{(l-m)!}{(l+m)!}}P_l^m(cos(\theta)e^{im\phi}), \tag{1}$$

where (θ,ϕ) and $\theta\in[0,\pi]$, $\phi\in[0,2\pi]$ correspond to the spherical coordinates, and P_l^m denotes the associated Legendre polynomials.

The modified SH functions [Tuc04] only consider even orders $\{k = 0,2,4,\dots,l\}$ with $m = \{-k,\dots,0,\dots,k\}$ to guarantee the required symmetry of the HARDI data:

$$Y_j = \begin{cases} \sqrt{2}\cdot Re(Y_k^m), & \text{if } -k\le m<0 \\ Y_k^0, & \text{if } m=0 \\ \sqrt{2}\cdot Img(Y_k^m), & \text{if } 0\le m<k \end{cases}, \tag{2}$$

where $Re(Y_k^m)$ and $Img(Y_k^m)$ denote the real and imaginary parts of Y_k^m respectively, and j represents the index in the modified basis, defined as $j = j(k,m) = (k^2+k+2)/2+m$.

Let $X\in\mathbb{R}^{d\times n}$ denote the matrix form for diffusion values from the raw HARDI data, where d is the number of diffusion gradient directions and $n = whs$ is the number of voxels. The matrix X is obtained by unfolding the HARDI data tensor along the fourth dimension. The SH representation, which is a linear transformation, can be given in matrix form as

$$X = SC = \begin{pmatrix} Y_1(\theta_1,\phi_1) & \dots & Y_r(\theta_1,\phi_1) \\ \vdots & \ddots & \vdots \\ Y_1(\theta_d,\phi_d) & \dots & Y_r(\theta_d,\phi_d) \end{pmatrix}\begin{pmatrix} c_{1,1} & \dots & c_{1,n} \\ \vdots & \ddots & \vdots \\ c_{r,1} & \dots & c_{r,n} \end{pmatrix}, \tag{3}$$

Hence, $S\in\mathbb{R}^{d\times r}$ contains the SH basis functions, with $r = (l+1)(l+2)/2$ is the number of modified SH basis functions, and $C\in\mathbb{R}^{r\times n}$ contains the SH coefficients. To directly solve (3) for C, we would have to invert S. However, since in practice $d > r$, the matrix S is not invertible. Descoteaux et al. [DAFD07] have shown that the linear system in (3) has the following closed-form solution:

$$C = \underbrace{(S^TS+\lambda L)^{-1}S^T}_{S^+}X = S^+X, \tag{4}$$

where λ is the weight on the regularization term and $L\in\mathbb{R}^{r\times r}$ is a diagonal matrix, called the Laplace-Beltrami smoothing matrix, with entries $l_j^2(l_j+1)^2$. Moreover, $l_j = \{0,2,\cdots,l\}$ is the order associated with the j_{th} coefficient, $j = \{1,2,\cdots,(l+1)(l+2)/2\}$.

3.2.2. Funk-Radon transform of modified SH

The Funk-Radon transform (FRT) of diffusion values X_i, i.e., a column in X, in a unit direction $\boldsymbol{u}$, is obtained by computing the integral over the great circle that is perpendicular to $\boldsymbol{u}$:

$$\int_{|\boldsymbol{w}|=1}\delta(\boldsymbol{u}^T\boldsymbol{w})X_i(\boldsymbol{w})\,d\boldsymbol{w}, \tag{5}$$

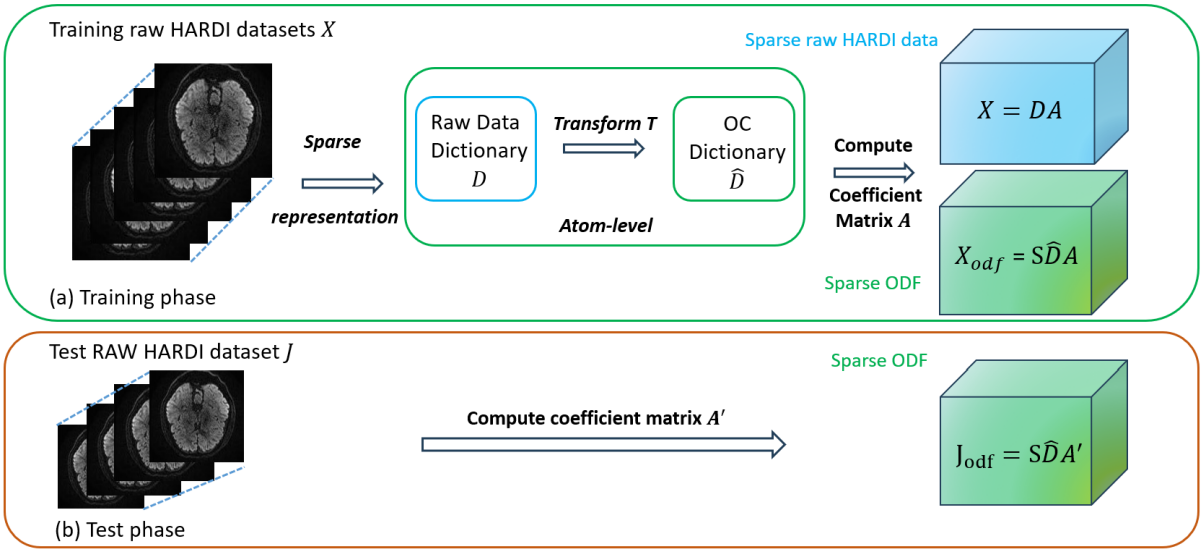

Figure 2: *Sparse pipeline for q-ball imaging: (a) illustrates the training phase, where we feed our pipeline with raw HARDI datasets and train a dictionary D. The dictionary D is a collection of basis functions (atoms), which capture representatives of the data. Then the corresponding dictionary for the ODFs $\hat{D}$ is computed on an atom-level using a transform T combining the spherical harmonics representation and funk-radon transform. By using OMP (orthogonal matching pursuit) to compute sparse coefficients α we obtain the sparse representations of the raw HARDI data and for the ODF. (b) exhibits the process of computing sparse ODF for new unseen raw HARDI data. With a trained dictionary from the training phase, one only needs to compute its sparse coefficients voxelwise to get sparse ODF.*

where $\delta(.)$ is the Dirac delta function. By replacing the signal X_i with its modified SH approximation, denoted H_l for degree l, the above equation is reformulated as

$$\int_{|\boldsymbol{w}|=1} \delta(\boldsymbol{u}^T \boldsymbol{w}) H_l(\boldsymbol{w})\, d\boldsymbol{w} = 2\pi P_l(0) H_l(\boldsymbol{u}), \tag{6}$$

where the Dirac delta function is replaced with a delta sequence $\delta_n(x) = (n/\sqrt{\pi}) exp(-n^2x^2)$ to ensure continuity in the interval $[-1,1]$. Moreover, $P_l(0)$ is the Legendre polynomial of degree l evaluated at 0. Hence, the Funk-Radon transform of a voxel X_i in the HARDI data on a unit vector direction $\boldsymbol{u}$ with the modified SH approximation is as follows:

$$G[X_i](\boldsymbol{u}) = \sum_{j=1}^{r} 2\pi P_{l_j}(0) C_{j,i} Y_j(\boldsymbol{u}). \tag{7}$$

Since we are typically interested in the SH coefficients of the ODF, we may rewrite (7) in matrix form as follows:

$$\hat{C} = PC = \begin{bmatrix} \ddots & & \\ & 2\pi(-1)^{\frac{l_j}{2}} \frac{1\cdot 3\cdot 5\cdot 7\cdots(l_j-1)}{2\cdot 4\cdot 6\cdot 8\cdots(l_j)} & \\ & & \ddots \end{bmatrix} C, \tag{8}$$

where the diagonal elements of P are $2\pi P_{l_j}(0)$, with $P_{l_j}(0)$ being the Legendre polynomial of degree l with nonzero values only when l is even and evaluated at 0. Fig 1 shows 2 examples of FRT.

4. Sparse Representation for raw HARDI data

A linear signal model, or representation, for the raw HARDI data $X \in \mathbb{R}^{d\times n}$ is expressed as $X = DA$, where $D \in \mathbb{R}^{d\times k}$ is called a *dictionary*, i.e., a matrix where each column is a basis vector, referred to as an *atom*. In essence, the dictionary, D, constructs a linear mapping from $\mathbb{R}^d$ to $\mathbb{R}^K$, and each column in the matrix A contains the representation coefficients. An example of such a representation is given in (3), where D contains SH basis functions and A is SH coefficients. When the matrix A is sparse, we achieve a sparse representation of X. To measure the sparsity of the representation, we use the ℓ_0 pseudo-norm, denoted $\|.\|_0$, which counts the number of non-zero elements in a vector. To achieve a sparse representation, we often require the dictionary to be *overcomplete*, i.e., $k > d$. Typically, the more overcomplete a dictionary, the sparser the representation. However, an overcomplete dictionary implies that solving the linear system $X = DA$ for A is underdetermined with infinitely many solutions. Therefore, we add a sparsity regularizer and solve the following sets of equations instead:

$$\min \ \|A_i\|_0 \quad s.t. \quad \|X_i - DA_i\|_2 \leq \varepsilon, \tag{9}$$

for all $i = \{1,\ldots,n\}$, i.e., for each voxel, where ε is a user-defined threshold for the representation error. Equation (9) can be solved efficiently using, e.g., Orthogonal Matching Pursuit (OMP) [PRK93].

It is well known that learning-based dictionaries achieve a higher level of sparsity in the representation as compared to analytical bases. The problem of dictionary learning is formulated as

$$\min_{\mathbf{D},\mathbf{A}} \|X - \mathbf{DA}\|_F^2 \quad s.t. \quad \|\mathbf{A}_i\|_0 \leq T_0, \ \forall i \in \{1,\ldots,n\}, \tag{10}$$

where T_0 is a user-defined sparsity parameter, and, with a slight abuse of notation, X is a HARDI training set. An efficient and commonly used algorithm for solving (10) is K-SVD [AEB06]. The K-SVD algorithm takes as input a training dataset, $\mathbf{X} \in \mathbb{R}^{d\times n}$, the number of atoms, k, and a sparsity parameter, T_0, in order to compute an overcomplete dictionary $\mathbf{D} \in \mathbb{R}^{d\times k}$. According to (10),

since the optimization problem has two unknown variables, we iteratively optimize for each variable until convergence. First, for each voxel in the training set, we compute a sparse representation with a maximum of T_0 nonzero coefficients. Next, we utilize the training set and the dictionary estimate from the previous iteration to update each atom of the dictionary in the current iteration. For the first iteration, the dictionary is typically initialized with random numbers. We refer the reader to [AEB06] for more implementation details.

4.1. Sparse representations and the q-ball imaging pipeline

In practice, the reconstruction of ODFs from HARDI data for often large patient groups is an expensive operation using the traditional numerical q-ball imaging pipeline. To overcome this limitation, we propose a novel pipeline shown in Figure 2 that combines q-ball imaging and sparse representation to estimate the ODF coefficients directly from the sparse representation of the raw HARDI data. Since our method relies on dictionary learning, we assume that the HARDI data is divided into a training and a testing set. The training set is used for learning an overcomplete dictionary, while the test set is used for utilizing the trained dictionary to obtain ODF coefficients. As a first step, we train an overcomplete dictionary D using a training set of raw HARDI data, denoted $X \in \mathbb{R}^{d\times n}$, with a user-defined sparsity parameter T_0. Note that the training set may contain voxels from multiple patients. Accordingly, we have

$$X_i = DA_i,\ \|A_i\|_0 \leq T_0,\ \forall i \in \{1,\dots,n\}. \tag{11}$$

The analytical q-ball imaging proposed by Descoteaux [DAFD07], as described in Section 3, computes SH coefficients of the ODF by evaluating

$$\hat{C} = PC = \underbrace{P(S^TS+\lambda L)^{-1}S^T}_{T} X = TX \tag{12}$$

A key observation in our proposed method is that one can transform the dictionary atoms using (12). This transforms the atoms of the original dictionary, trained on raw data, to reside in the ODF coefficient space. As such, we can write:

$$\hat{D} = P(S^TS+\lambda L)^{-1}S^TD = TD, \tag{13}$$

where each atom of $\hat{D}$ contains ODF coefficients of the corresponding atom in D.

Let $J \in \mathbb{R}^{d\times m}$ be a HARDI data for one patient in the test set. Using (12) and the trained dictionary on raw HARDI data, we have $\hat{C} = TJ = TDA'$, where, with a slight abuse of notation, $\hat{C}$ now denotes the ODF coefficients of the test set. Moreover, A' is the sparse coefficient matrix of the test set. Hence, to obtain A', we need to solve

$$\min \|A'_i\|_0 \ \ s.t.\ \ \|J_i - DA'_i\|_2 \leq \varepsilon \tag{14}$$

for each voxel J_i, $i = \{1,\dots,m\}$ to obtain A. An important outcome of (14) is that each voxel in J will have an adaptive sparsity defined over the error threshold ε. Such per-voxel sparsity can be utilized for visualizing the volume, where low sparsity voxels (i.e., a high number of nonzero coefficients) will mark regions of interest.

Once the sparse coefficients A' are recovered using (14), the ODF coefficients of Y are computed efficiently via

$$\hat{C} = TDA' = \hat{D}A', \tag{15}$$

and therefore the ODF of J is $J_{odf} = S\hat{D}A'$. Since $S\hat{D}$ can be precomputed in the training phase, obtaining the ODF becomes a multiplication of a dense matrix by a sparse matrix, which can be computed efficiently using a GPU or CPU implementation. Most importantly, we only have to apply the q-ball imaging pipeline, i.e., the transformation T, on the atoms of the dictionary and not for every voxel of the HARDI data, which reduces the computational cost significantly. As an example, a dictionary has typically 10^2 atoms, while the HARDI data contains 10^7 voxels. The computational complexity of Analytical QBI is $\mathcal{O}(whsdr)$, while for our pipeline is $\mathcal{O}(whs\beta r)$, where β is the average sparsity of voxels. Since $\beta < d$, and in most cases $\beta \ll d$, our method is significantly faster than traditional QBI. Another important aspect of our method is that the dictionary can be computed once and then reused for each any HARDI dataset. This leads to a significantly reduced storage cost, since only the sparse coefficients A are required for new HARDI datasets; see Section 5. Finally, according to (15), we observe that the sparse coefficient matrix A' is shared between the raw HARDI and ODF coefficient fields. This implies that an analysis of the sparse coefficients, e.g., for fiber tracking, holds for both fields. We have left such in-depth analysis to future work.

5. Experimental results and visualization

This section presents an evaluation of sparse modelling of HARDI data. First, we explore the distribution of diffusion values, both visually and statistically, by comparing different regions within the same volume. We study the choice of parameters and evaluate the performance of our pipeline using both efficiency-related metrics evaluating the computational complexity and quality-related metrics evaluating the visual performance. Finally, we present a visual evaluation of the dictionaries in ODF space exploring how the sparse coefficients are dynamically distributed for different data depending on voxel complexity.

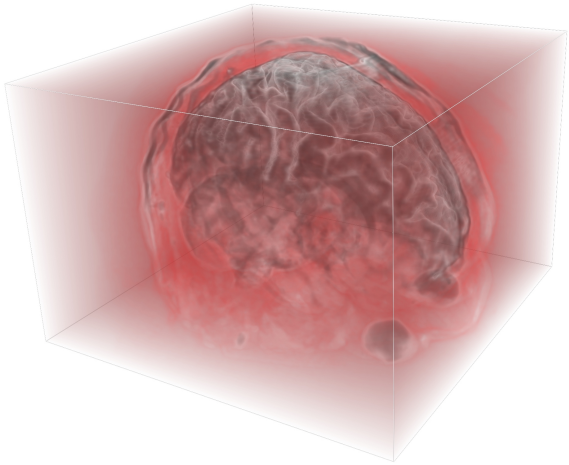

Figure 3: *Volume rendering of the brain showing one of the 64 diffusion values. The red color represents low diffusion values mostly outside the brain, while the higher values are shown in gray highlighting regions within the brain. Inviwo [JSS*19].*

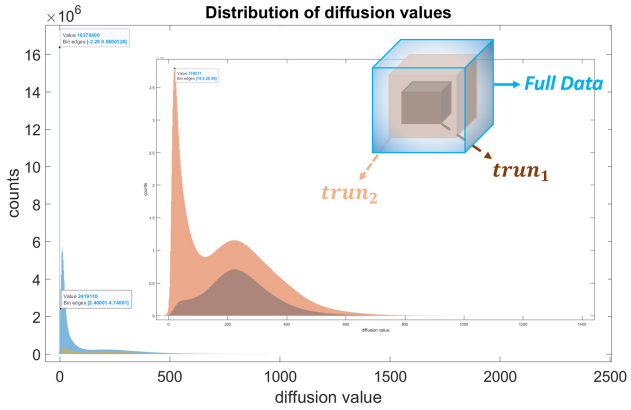

We used

Figure 4: *Comparisons of the distribution of diffusion values over the original data of size* $140 \times 140 \times 90 \times 64$*,* $trun_1$ *of size* $70 \times 70 \times 50 \times 64$*, and* $trun_2$ *of size* $100 \times 100 \times 60 \times 64$*. The original data has a large peak near the zero-value, which can be associated with out-of-bound voxels. The close-up only shows the distribution of the two cropped datasets. As* $trun_1$ *does not include any out-of-bound voxels the histogram does not show a peak around zero.*

5.1. Exploration over diffusion values

To get a full picture of the performance of our proposed pipeline we decided not to use a mask for the brain but use the full raw dataset. We start by exploring the data visually and build an overview of the distribution of the diffusion values. The visualizations explore the 64 diffusion values as well as the magnitude of the signal corresponding to the different directions, see Figure 3 for an example. As expected, there is noise present outside the brain with generally quite low diffusion values, however, there is also a region around the brain with extremely high values. To get a better understanding of the influence of noise on the resulting dictionary and the reconstruction quality we constructed two truncated datasets with different sizes by removing the outer parts of the dataset as shown in Figure 4. We label the resulting volumes as $trun_1$ and $trun_2$. The histogram of diffusion values from the original data, $trun_1$, and $trun_2$ are also shown in Figure 4 which also shows a close-up for $trun_1$, and $trun_2$ since the dynamic range of the original data is very large dominating the histogram. The maximum value of raw data for full data, $trun_1$, $trun_2$ is about 2000, 1200, and 1500, respectively. The maximum value of ODF for full data, $trun_1$, $trun_2$ is around 4000, 2000, and 3800, respectively.

5.2. Performance Evaluation of our pipeline

We evaluate our pipeline with respect to efficiency, in Section 5.2.1, and quality, in Section 5.2.2. For efficiency, we utilize two metrics. The first is the total time it takes for the pipeline to complete the task of obtaining ODF coefficients from raw data. The second metric is compression ratio, which is the ratio between the original data size and the size of the compressed data, which in our case corresponds to the sparse coefficients. The compression ratio is an important metric for data transmission and storage. For quality, we use Root Mean Square Error (RMSE), which measures the average differences between original data and reconstructed data after sparse representation is applied.

Figure 5: *Illustration for training and testing dataset choices for cross-validation. Blue represents training data, and orange stands for test data for this group. For each run, we take 4 volumes for training and 1 volume for testing. The five different train/test divisions are labeled as No. 1, No. 2, No. 3, No. 4, and No. 5.*

We run two types of experiments to test our pipeline. The first type of experiment, denoted as $Exp : Individual$ means the training and test datasets are identical. The second type of experiment, denoted as $Exp : Group$ is to train and test from different datasets to see how robust the pipeline will be when fed with new unseen data. And here we choose to use the cross-validation method for our available five different pre-processed datasets, as shown in Figure 5, where the pre-processing involves motion correction for gradient directions. Cross-validation is commonly used in machine learning to evaluate and compare different models over new unseen datasets by using different portions of data to train and test.

5.2.1. Efficiency

In Table 1, we report timing results for our pipeline, taking into account the original data, $trun_1$, and $trun_2$. The last column represents the main timing results, where the total computation time includes obtaining sparse coefficients using the OMP algorithm [PRK93], i.e. solving (14), as well as computing the ODF coefficients, i.e. by evaluating (15). We would like to emphasize that since the training is a one-time process, it cannot be included in the total computation time of our QBI pipeline. However, for completeness, the table also includes the total computation time for training.

Figure 6 shows the compression ratio with respect to RMSE for $Exp : Individual$ and $Exp : Group$. We compare results in both the raw data field and the ODF field. The compression ratio is defined as the ratio between the size of uncompressed data and compressed data, e.g. a compression ratio of 4 means the compressed data is 4 times smaller. Specifically, since we are compressing the diffusion values using sparse coefficients, it is the ratio between the total number of diffusion values over the total number of nonzero coefficients of the sparse representation. Moreover, because the size of the dictionary D and the transformation matrix T are negligible with respect to the size of the HARDI data, we have not included them in the compression ratio calculation.

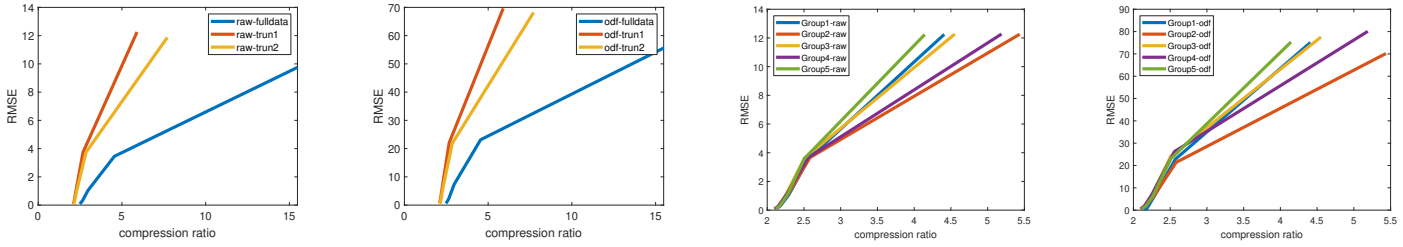

(a) RMSE of our method w.r.t compression ratio for raw data and *Exp* : *Individual*

(b) RMSE of our method w.r.t compression ratio for ODF and *Exp* : *Individual*

(c) RMSE of our method w.r.t compression ratio for raw data and *Exp* : *Group*

(d) RMSE of our method w.r.t compression ratio for ODF and *Exp* : *Group*

Figure 6: *Depending on the chosen accuracy by the RMSE, different compression ratios can be achieved. For the minimal RMSE close to zero, the compression rate is about 2 in all cases. For a fixed RMSE, we observe that different datasets show distinct compression ratios, implying that their coefficients have different sparsity levels. The data range for different data sets is presented in Section 5.1. For (c) and (d) we use* $trun_1$.

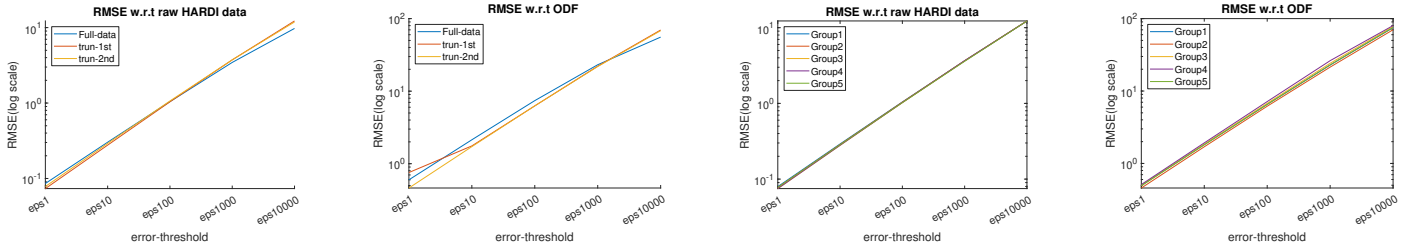

(a) RMSE of our method w.r.t ε for raw data and *Exp* : *Individual*

(b) RMSE of our method w.r.t ε for ODF and *Exp* : *Individual*

(c) RMSE of our method w.r.t ε for raw data and *Exp* : *Group*

(d) RMSE of our method w.r.t ε for ODF and *Exp* : *Group*

Figure 7: *Comparisons of RMSE with respect to the error threshold* ε *in* (14) *for raw HARDI data and their ODF. For (c) and (d) we use* $trun_1$.

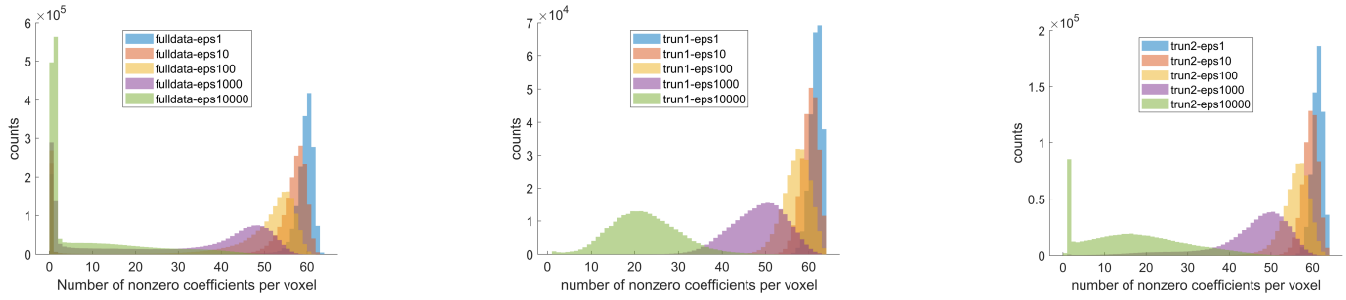

(a) Histogram of nonzero coefficients for full data

(b) Histogram of nonzero coefficients for $trun_1$

(c) Histogram of nonzero coefficients for $trun_2$

Figure 8: *Exp* : *Individual. Comparisons of sparsity level of coefficients under different error threshold levels for expressing full and truncated data. Both full and* $trun_2$ *have a peak for zero coefficients and show higher sparsity under eps10000.*

5.2.2. Quality

Root Mean Square Error (RMSE) is a common evaluation metric for evaluating the quality of reconstructed data by investigating how much the reconstruction deviates from the original data from the perspective of the standard deviation. It measures the difference globally between the original and the reconstructed data. A smaller reconstruction error indicates a better reconstruction. We compare reconstructions of raw HARDI data and ODF over three different data sizes (full data, $trun_1$, and $trun_2$) with respect to the error threshold ε in (14). We label the experiments as e.g. eps10, implying that $\varepsilon = 10$. Figure 7 demonstrates our pipeline's RMSE for a variable threshold ε. Indeed, as the error threshold for the sparse reconstruction increases, the RMSE also increases. This pattern is evident for raw and ODF data, as well as different cross-validation sets, although the rate of change in RMSE with respect to ε varies between different experiments. Figure 13 illustrates a reconstructed raw HARDI slice followed by an error distribution map computed from the original slice, where it can be seen that the error is uniformly distributed regardless of the error threshold, ε. In Figure 8, we plot the histogram of sparsity, i.e. the number of nonzero coefficients in A'. As expected, for a small ε, the majority of voxels use a large number of coefficients, hence the represetation is less sparse.

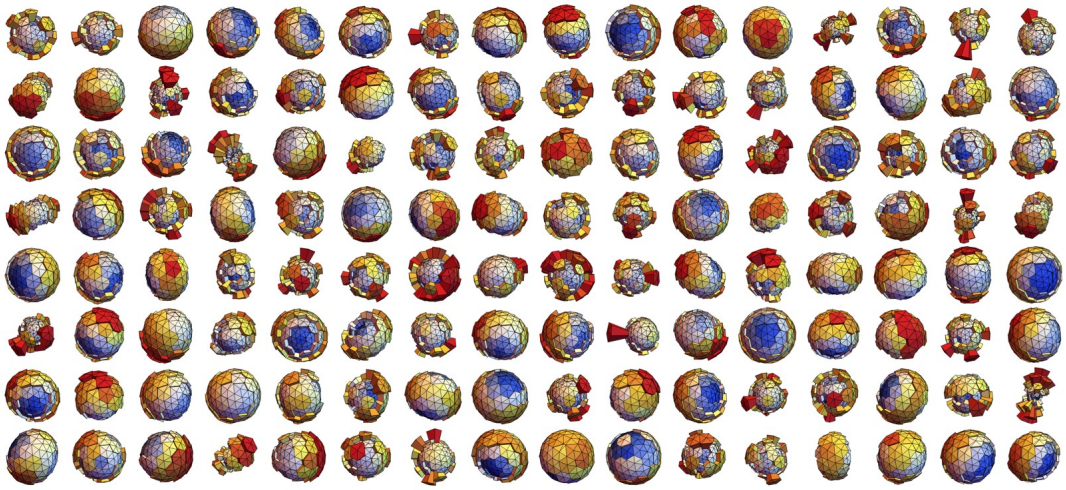

Figure 9: *Illustrates a dictionary for raw HARDI data trained with full data. The glyph uses a Voronoi diagram of the sphere, assigning regions to every direction. The length and color of each region indicate the magnitude of the diffusion direction. Blue indicates a low value, and red suggests a high value.*

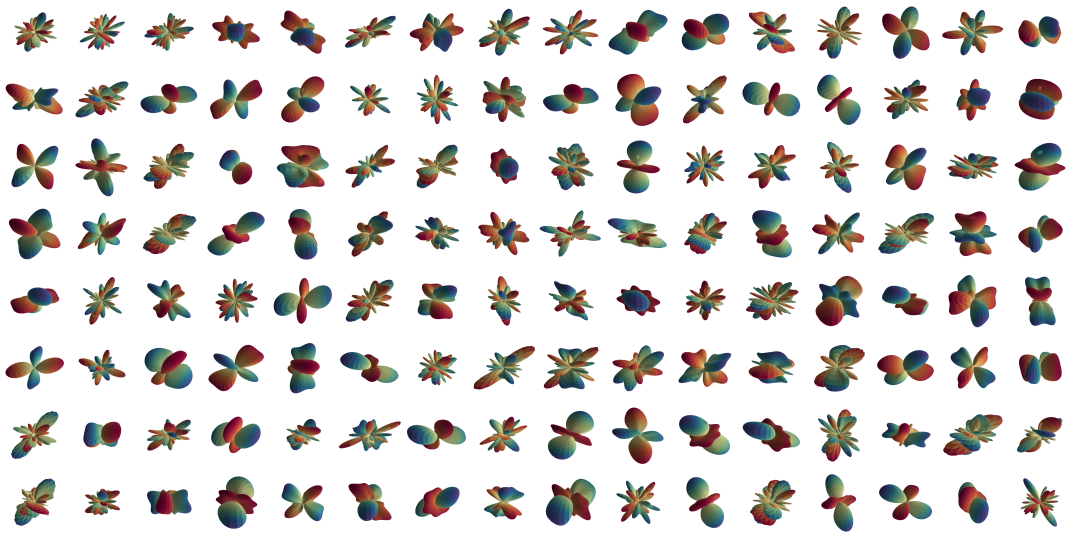

Figure 10: *Illustrates a dictionary for ODF trained from Group 2 with $trun_1$ since it produces the lowest RMSE among 5 groups. These per-atom glyphs represent projected ODF coefficients on spherical harmonics bases, i.e. $STD = S\hat{D}$, hence they are symmetric.*

5.3. Visual inspection of HARDI dictionaries

As a first step towards evaluating the dictionaries, we provide visualizations of the atoms of the dictionary. Glyph representation is commonly used in medical and scientific data visualization to represent multi-dimensional data. While there are glyphs for HARDI data [SV19], we first introduce a glyph that directly renders the raw data without applying any interpretation. The glyphs highlight all 64 diffusion directions. Each direction is represented by a corresponding Voronoi cell on the sphere. The radius is scaled by the diffusion value of this channel. The color contains redundant information. This representation does not respect the symmetry of the data with respect to diffusion directions. Figure 9 shows the dictionary D trained on raw data with 128 atoms and using the same glyph representations. Each glyph represents one atom in the dictionary for the full dataset. Figure 10 shows the atoms of the dictionary D as ODFs using spherical harmonics from Group 2 with $trun_1$. In other words, we use glyphs to visualize $STD = S\hat{D}$. Here, as expected, the atoms demonstrate symmetry due to the spherical harmonics representation.

In the next step, we are interested in how frequently the atoms are used in the dataset. In the sparse representation literature, see e.g. [Ela10], the index of atoms used for a representation is called the *support*. The histogram of the support taken over all the voxels in the data can reveal the importance of atoms in the dictionary. Figure 11(a-c) demonstrates the histogram of support for $Exp : Group$ for three distinct values of ε. Figure 11(d-f) shows the results for $Exp : Individual$ under a similar setup. Interestingly, in both cases,

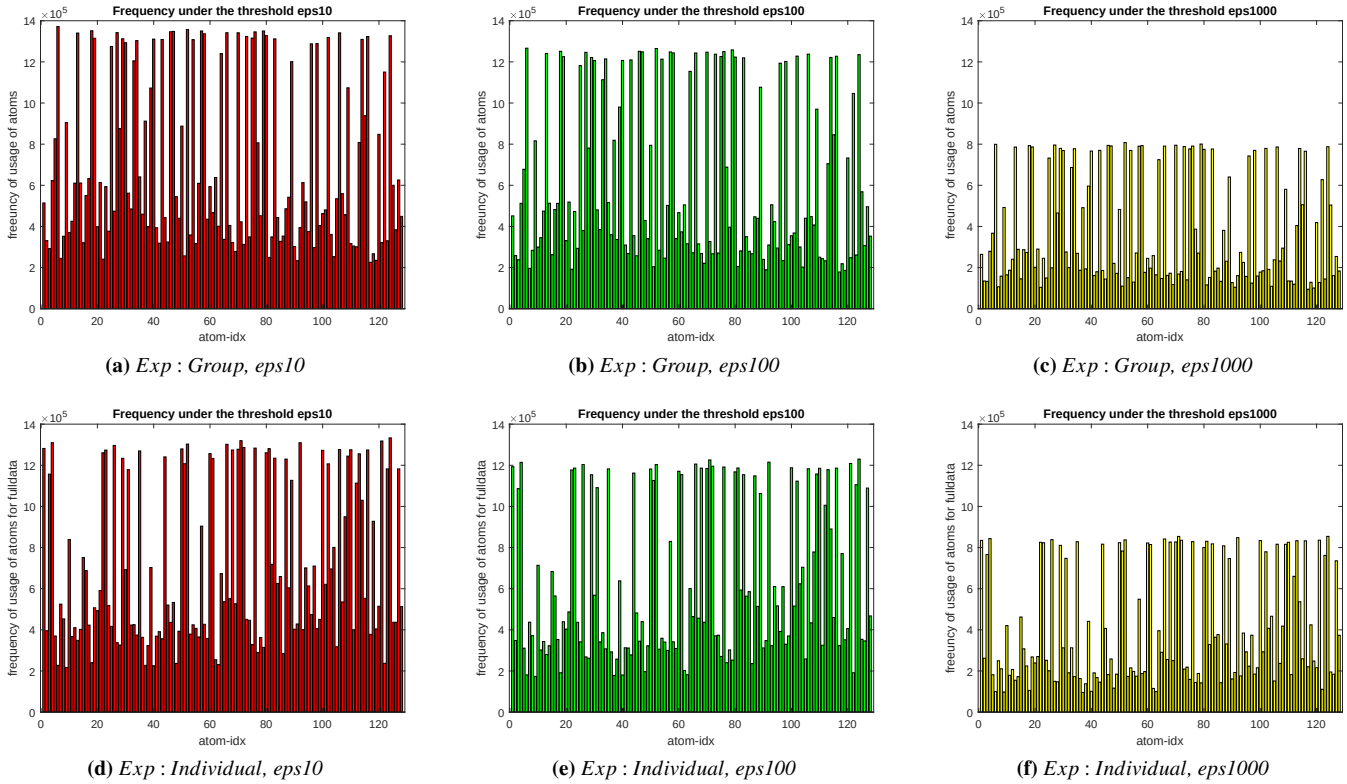

(a) *Exp* : *Group, eps10* **(b)** *Exp* : *Group, eps100* **(c)** *Exp* : *Group, eps1000*

(d) *Exp* : *Individual, eps10* **(e)** *Exp* : *Individual, eps100* **(f)** *Exp* : *Individual, eps1000*

Figure 11: *Histograms of the frequency of atoms using different error levels for reconstructing the raw HARDI data.*

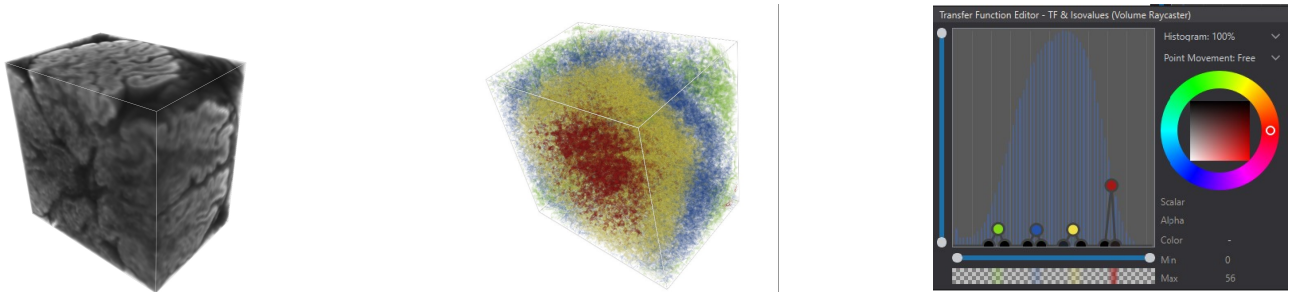

(a) *raw data* **(b)** *sparsity map* **(c)** *transfer function*

Figure 12: *Exp* : *Individual. illustrate the sparsity of voxels in* $trun_1$ *region. The sparse coefficients follow a symmetric Gaussian distribution. We use colors to represent different sparsity levels, e.g., green represents the region with the most sparse coefficients followed by blue, yellow, and red stands for the region that requires the densest coefficients.*

we observe that by changing the threshold, the importance of the atoms, demonstrated by their frequency of usage, does not change. This implies that the importance of atoms in the sparse representation does not vary with respect to the desired reconstruction quality. We also observe, in both figures, that the frequency becomes lower when the error threshold is increased, which is expected since sparsity increases and fewer atoms contribute to the representation. Apart from the support, another informative quantity is sparsity, i.e. the number of atoms used in the representation. In Figure 12, we visualize the distribution of the sparsity of HARDI voxels, highlighting four separate regions of sparsity in the volume.

6. Discussion and future work

This paper demonstrated a novel sparse signal model for visual exploration and analysis of diffusion MRI HARDI data. We showed that sparse modeling can be efficiently integrated into the q-ball imaging pipeline, yielding high-performance and high-quality results. Sparse representations lead to significant reductions in the amount of data required to represent the HARDI data. The compression is not lossless, however, even with very small error thresholds already good compression ratios can be achieved. Our evaluation and results showed that the computational complexity of the analytical QBI of our pipeline in comparison to the standard ap-

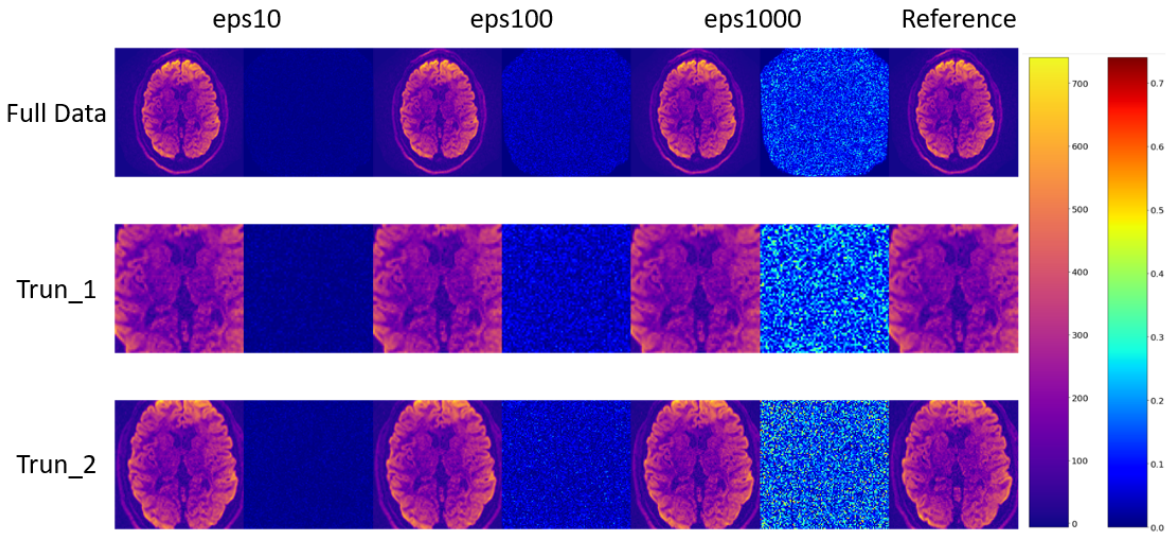

Figure 13: *Error distribution maps of a slice of the reconstructed HARDI data under different threshold levels. The left image for each threshold value is the reconstructed slice and the right image is the error distribution map. Left color bar indicates the range of values for signal magnitude, and the right color bar shows relative error normalized using the maximum error of eps1000.*

Dataset	Size($w \times h \times s \times d$)	Time (sec) for training dictionary	Time (sec) for computing Coefficients (eps1)	Time (sec) for computing Coefficients(eps100)	Time (sec) for computing Coefficients (eps10000)
Full data	$140 \times 140 \times 96 \times 64$	413.39	17.17	14.58	3.52
$trun_1$	$70 \times 70 \times 50 \times 64$	57.68	2.72	2.63	0.94
$trun_2$	$100 \times 100 \times 60 \times 64$	150.94	6.46	6.04	1.82
Group 1-full data	$140 \times 140 \times 384 \times 64$	1619.56	17.92	14.88	3.33
Group 1- $trun_1$	$70 \times 70 \times 200 \times 64$	256.78	2.66	2.41	1.25
Group 2-$trun_1$	$70 \times 70 \times 200 \times 64$	260.76	2.80	2.53	1.02
Group 3-$trun_1$	$70 \times 70 \times 200 \times 64$	245.66	2.74	2.47	1.21
Group 4-$trun_1$	$70 \times 70 \times 200 \times 64$	250.47	2.91	2.44	1.08
Group 5-$trun_1$	$70 \times 70 \times 200 \times 64$	252.57	2.83	2.52	1.31

Table 1: *Comparisons of computing time for dictionary and sparse coefficients over different datasets under different error threshold levels. All experiments are run with SH basis of order $l = 8$. Computation is performed on a 16-core processor at 3.40 GHz.*

proach is β/d, where β is the average sparsity of voxels and d is the number of measured diffusion directions. Since $\beta < d$, in most cases $\beta << d$, our approach is significantly faster than traditional QBI. An additional advantage is that the sparse coefficient matrix A' is shared between the raw HARDI and ODF coefficient fields.

The promising results from our experiments show that there is great value in exploring sparse modelling of HARDI data further. There are several venues for future work. For example, in our evaluation, we used an overcomplete K-SVD dictionary. There are a wide variety of dictionary formulations with different pros. and cons., and we believe that it would be interesting to further investigate different choices, especially in the context of data analysis directly from, e.g., clustering in coefficient space. Looking at the dictionary we can see that many atoms appear in a similar form in a similar orientation. This behaviour can be avoided by enforcing rotationally invariant atoms. This would also lead to an even more compact representation and open up for novel data analysis. A trained dictionary and its sparse coefficients benefit both data transmission and storage. A long-term goal is to exploit sparse representations and data-specific dictionaries for the extraction of patterns usable in visualization.

Acknowledgment

This work is supported by the Excellence Center at Linköping and Lund in Information Technology (ELLIIT) for strategic research in Sweden, SeRC (Swedish e-Science Research Center), Wallenberg Autonomous Systems and Software Program (WASP), and the Swedish Research Council (VR) grant 2019-05487. We would like to thank Peter Steneteg and Martin Falk for their support with the visualizations.

References

[AAB*15] AMYOT F., ARCINIEGAS D. B., BRAZAITIS M. P., CURLEY K. C., DIAZ-ARRASTIA R., GANDJBAKHCHE A., HERSCOVITCH P., HINDS S. R., MANLEY G. T., PACIFICO A., ET AL.: A review of the effectiveness of neuroimaging modalities for the detection of traumatic brain injury. *Journal of neurotrauma 32*, 22 (2015). 2

[AEB06] AHARON M., ELAD M., BRUCKSTEIN A.: K-svd: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on signal processing 54*, 11 (2006), 4311–4322. 4, 5

[AJBTdS19] ASSAF Y., JOHANSEN-BERG H., THIEBAUT DE SCHOTTEN M.: The role of Diffusion MRI in neuroscience. *NMR in Biomedicine - Special Issue: Diffusion MRI of the brain: The naked truth 32*, 4 (2019). 2

[AN10] ASHOK A., NEIFELD M. A.: Compressive Light Field Imaging. vol. 7690, pp. 1–12. 2

[BML94] BASSER P. J., MATTIELLO J., LEBIHAN D.: Mr diffusion tensor spectroscopy and imaging. *Biophysical journal 66*, 1 (1994), 259–267. 1, 2

[BRGIG*14] BALSA RODRÍGUEZ M., GOBBETTI E., IGLESIAS GUITIAN J. A., MAKHINYA M., MARTON F., PAJAROLA R., SUTER S. K.: State-of-the-art in compressed gpu-based direct volume rendering. In *Computer graphics forum* (2014), vol. 33, Wiley Online Library, pp. 77–100. 2

[CCFa] CCF: Human connectome project. https://www.humanconnectome.org/study/hcp-young-adult. Accessed: 2023-01-06. 3

[CCFb] CCF: Mgh adult diffusion data acquisition details. https://www.humanconnectome.org/study/hcp-young-adult/document/mgh-adult-diffusion-data-acquisition-details/. Accessed:23-01-06. 3

[CEX88] CALLAGHAN P. T., ECCLES C. D., XIA Y.: Nmr microscopy of dynamic displacements: k-space and q-space imaging. *J. Phys. E: Sci Instrum 21* (1988), 820–822. 1

[CT20] CHRISTIAENS D., TOURNIER J. D.: *Advances in Magnetic Resonance Technology and Applications*. Elsevier, 2020, ch. 20 Modeling Fiber Orientations Using Diffusion MRI, pp. 509–532. 1

[DAFD07] DESCOTEAUX M., ANGELINO E., FITZGIBBONS S., DERICHE R.:. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine 58*, 3 (2007), 497–510. 3, 5

[DDT*08] DUARTE M. F., DAVENPORT M. A., TAKHAR D., LASKA J. N., SUN T., KELLY K. F., BARANIUK R. G.: Single-pixel Imaging via Compressive Sampling. *IEEE Signal Processing Magazine 25*, 2 (March 2008), 83–91. 2

[Des15] DESCOTEAUX M.: *Wiley Encyclopedia of Electrical and Electronics Engineering*. Wiley, 2015, ch. High Angular Resolution Diffusion Imaging (HARDI). 2

[DMG20] DÍAZ J., MARTON F., GOBBETTI E.: Interactive spatiotemporal exploration of massive time-varying rectilinear scalar volumes based on a variable bit-rate sparse representation over learned dictionaries. *Computers & Graphics 88* (2020), 45–56. 2

[Ela10] ELAD M.: *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing*, 1st ed. Springer Publishing Company, Incorporated, 2010. 8

[F*16] FAN Q., ET AL.: Mgh–usc human connectome project datasets with ultra-high b-value Diffusion MRI. *Neuroimage 124* (2016), 1108–1114. 3

[GIGM12] GOBBETTI E., IGLESIAS GUITIÁN J. A., MARTON F.: Covra: A compression-domain output-sensitive volume rendering architecture based on a sparse representation of voxel blocks. In *Computer Graphics Forum* (2012), vol. 31, Wiley Online Library, pp. 1315–1324. 2

[GvdVS22] GRUEN J., VAN DER VOORT G., SCHULTZ T.: Model averaging and bootstrap consensus-based uncertainty reduction in diffusion mri tractography. *Computer Graphics Forum xx*, xx (2022). 2

[HKDMS17] HEIM B., KRISMER F., DE MARZI R., SEPPI K.: Magnetic resonance imaging for the diagnosis of parkinson's disease. *Journal of neural transmission 124*, 8 (2017), 915–964. 2

[HWG*21] HAN S., WANG N., GUO Y., TANG F., XU L., JU Y., SHI L.: Application of sparse representation in bioinformatics. *Frontiers in Genetics 12* (2021). 1

[JLD13] JIANG Z., LIN Z., DAVIS L. S.: Label Consistent K-SVD: Learning a Discriminative Dictionary for Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence 35*, 11 (Nov 2013), 2651–2664. 2

[JPGJ12] JIAO F., PHILLIPS J. M., GUR Y., JOHNSON C. R.: Uncertainty visualization in HARDI based on ensembles of odfs. In *2012 IEEE Pacific Visualization Symposium* (2012), IEEE, pp. 193–200. 1

[JS23] JUMAKULYYEV I., SCHULTZ T.: Combining image space and q-space pdes for lossless compression of diffusion mr images. *Journal of Mathematical Imaging and Vision* (2023). 2

[JSS*19] JÖNSSON D., STENETEG P., SUNDÉN E., ENGLUND R., KOTTRAVEL S., FALK M., YNNERMAN A., HOTZ I., ROPINSKI T.: Inviwo—a visualization system with usage abstraction levels. *IEEE transactions on visualization and computer graphics 26*, 11 (2019), 3241–3254. 5

[KVJ*21] KARIMI D., VASUNG L., JAIMES C., MACHADO-RIVAS F., KHAN S., WARFIELD S. K., GHOLIPOUR A.: A machine learning-based method for estimating the number and orientations of major fascicles in diffusion-weighted magnetic resonance imaging. *Med Image Anal 72* (2021), 102129. 2

[LDP07] LUSTIG M., DONOHO D., PAULY J. M.: Sparse MRI: The Application of Compressed Sensing for Rapid MR Imaging. *Magnetic Resonance in Medicine 58*, 6 (2007), 1182–1195. 2

[Lee19] LEEMANS A.: *Diffusion MRI of the brain: The naked truth*. NMR Biomed, 2019. 1

[MHGP*21] MARTINEZ-HERAS E., GRUSSU F., PRADOS F., SOLANA E., LLUFRIU S.: Diffusion-weighted imaging: Recent advances and applications. In *Seminars in Ultrasound, CT and MRI* (2021), Elsevier, pp. 490–506. 1

[MT14] MORI S., TOURNIER J.: *Introduction to Diffusion Tensor Imaging and Higher Order Models*. Elsevier, 2014. 3

[NDH*15] NEHER P. F., DESCOTEAUX M., HOUDE J.-C., STIELTJES B., MAIER-HEIN K. H.: Strengths and weaknesses of state of the art fiber tractography pipelines – a comprehensive in-vivo and phantom evaluation study using tractometer. *Medical Image Analysis 26* (2015), 287–305. 2

[PPvA*09] PEETERS T., PRCKOVSKA V., VAN ALMSICK M., VILANOVA A., TER HAAR ROMENY B.: Fast and sleek glyph rendering for interactive HARDI data exploration. In *Proceedings of IEEE Pacific Visualization Symposium* (2009). 2

[PRK93] PATI Y., REZAIIFAR R., KRISHNAPRASAD P.: Orthogonal matching pursuit: recursive function approximation with applications to wavelet decomposition. In *Proceedings of 27th Asilomar Conference on Signals, Systems and Computers* (1993), vol. 1, pp. 40–44. 4, 6

[RB11] RAVISHANKAR S., BRESLER Y.: MR Image Reconstruction From Highly Undersampled k-Space Data by Dictionary Learning. *IEEE Transactions on Medical Imaging 30*, 5 (May 2011), 1028–1041. 2

[RMS*11] RATHI Y., MICHAILOVICH O., SETSOMPOP K., BOUIX S., SHENTON M., WESTIN C.-F.: Sparse multi-shell diffusion imaging. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2011* (2011). 1

[ST65] STEJSKAL E., TANNER J.: Spin diffusion measurements: Spin echoes in the presence of a time dependent field gradient. *J. Chem Phys 42*, 288 (1965). 3

[SV19] SCHULTZ T., VILANOVA A.: Diffusion MRI visualization. *NMR in Biomedicine - Special Issue: Diffusion MRI of the brain: The naked truth 32*, 4 (2019). 1, 2, 8

[SVC18] SCHWAB E., VIDAL R., CHARON N.: Joint spatial-angular sparse coding for dmri with separable dictionaries. *Medical Image Analysis 48* (2018), 25–42. 2

[SW03] SCHNEIDER J., WESTERMANN R.: Compression domain volume rendering. In *IEEE Visualization, 2003. VIS 2003.* (2003), IEEE, pp. 293–300. 2

[SWK10] SCHULTZ T., WESTIN C.-F., KINDLMANN G.: Multi-diffusion-tensor fitting via spherical deconvolution: A unifying framework. *Med Image Comput Comput Assist Interv 13*, 1 (2010), 674–681. 2

[SZ19] SOTIROPOULOS S. N., ZALESKY A.: Building connectomes using Diffusion MRI: why, how and but. *NMR in Biomedicine - Diffusion MRI of the brain: The naked truth 32*, 4 (2019). 2

[Tuc04] TUCH D. S.: Q-ball imaging. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine 52*, 6 (2004), 1358–1372. 1, 2, 3

[VRG22] VAISH A., RAJWADE A., GUPTA A.: Tl-HARDI: Transform learning based accelerated reconstruction of HARDI data. *Computers in Biology and Medicine 143* (2022), 105212. 2

[WMM*09] WRIGHT J., MA Y., MAIRAL J., SAPIRO G., HUANG T.: Sparse representation for computer vision and pattern recognition. *Proceedings of the IEEE 98*, 6 (2009), 1031–1044. 1

[YSL13] YAN R., SHAO L., LIU Y.: Nonlocal Hierarchical Dictionary Learning Using Wavelets for Image Denoising. *IEEE Transactions on Image Processing 22*, 12 (Dec 2013), 4689–4698. 2

[ZCH*17] ZHANG C., CAAN M. W., HÖLLT T., EISEMANN E., VILANOVA A.: Overview+ detail visualization for ensembles of diffusion tensors. *Computer Graphics Forum 36*, 3 (2017), 121–132. 2

[ZXY*15] ZHANG Z., XU Y., YANG J., LI X., ZHANG D.: A survey of sparse representation: algorithms and applications. *IEEE access 3* (2015), 490–530. 1