



# Combining Image Space and q-Space PDEs for Lossless Compression of Diffusion MR Images

Ikram Jumakulyyev<sup>1</sup> · Thomas Schultz<sup>1</sup>

Received: 14 June 2022 / Accepted: 2 March 2023 / Published online: 28 March 2023  
© The Author(s) 2023

## Abstract

Diffusion MRI is a modern neuroimaging modality with a unique ability to acquire microstructural information by measuring water self-diffusion at the voxel level. However, it generates huge amounts of data, resulting from a large number of repeated 3D scans. Each volume samples a location in q-space, indicating the direction and strength of a diffusion sensitizing gradient during the measurement. This captures detailed information about the self-diffusion and the tissue microstructure that restricts it. Lossless compression with GZIP is widely used to reduce the memory requirements. We introduce a novel lossless codec for diffusion MRI data. It reduces file sizes by more than 30% compared to GZIP and also beats lossless codecs from the JPEG family. Our codec builds on recent work on lossless PDE-based compression of 3D medical images, but additionally exploits smoothness in q-space. We demonstrate that, compared to using only image space PDEs, q-space PDEs further improve compression rates. Moreover, implementing them with finite element methods and a custom acceleration significantly reduces computational expense. Finally, we show that our codec clearly benefits from integrating subject motion correction and slightly from optimizing the order in which the 3D volumes are coded.

**Keywords** PDE-based image compression · DMRI · Q-space inpainting · Lossless medical image compression

## 1 Introduction

With the development of new medical imaging techniques and constant refinement of existing ones, the associated storage requirements have been reported to grow exponentially each year [1]. This explains why medical image compression is an active area of research.

Our work belongs to the family of compression algorithms that are based on partial differential equations (PDEs). The general idea behind this approach is to store a sparse subset of the image information and to reconstruct the remaining image via PDE-based inpainting [2].

PDE-based compression has a long tradition for the lossy compression of natural images [2,3] and videos [4–6]. The benefit of PDE-based approaches relative to transform-based codecs like JPEG [7] and JPEG2000 [8] has often been most

pronounced at high compression rates [3]. Even though this strategy for lossy compression has also been transferred to three-dimensional images [9], in medical imaging, lossless compression is often preferred to ensure that all diagnostically relevant details are preserved. In some cases, it is even legally forbidden to apply lossy compression for medical image archival [10,11].

We recently introduced a PDE-based codec for 3D medical images that stores the residuals between the original image and an intermediate PDE-based reconstruction to ensure that the final reconstruction is lossless, and we demonstrated that this strategy led to competitive compression rates [12]. In our current work, we extend this idea for the specific use case of image datasets from diffusion MRI.

Diffusion MRI (dMRI) [13,14] is a variant of magnetic resonance imaging in which diffusion sensitizing gradients are introduced into the measurement sequence. If the hydrogen nuclei that generate the MR signal undergo a net displacement along the gradient direction during the measurement, the signal is attenuated. Assuming that these displacements result from (self-) diffusion, comparing diffusion-weighted to non-weighted measurements permits computation of an apparent diffusion coefficient.

✉ Thomas Schultz  
schultz@cs.uni-bonn.de

Ikram Jumakulyyev  
s6ikjuma@uni-bonn.de

<sup>1</sup> B-IT and Department of Computer Science II, University of Bonn, Friedrich-Hirzebruch-Allee 5, 53115 Bonn, Germany

Taking measurements with different gradient directions captures the directional dependence of the diffusivity. It results from interactions between water and tissue microstructure and therefore carries information about structures that are much smaller than the MR image resolution. Important applications of dMRI include the detection of microstructural changes that are related to aging or disease, and the reconstruction of major white matter tracts, which is referred to as fiber tracking or tractography [15].

The large number of repeated measurements in diffusion MRI leads to large amounts of data. In practice, resulting image datasets are often compressed using GZIP [16]. In our previous work [12], we demonstrated that, compared to this, PDE-based lossless compression can further reduce the memory requirement of individual dMRI volumes by more than 25%. However, applying our codec to each 3D volume independently does not exploit the fact that measurements for nearby gradient directions are usually similar. Moreover, it is relatively time consuming.

In our current work, we address both of these limitations by combining the previous idea of lossless compression via image-space inpainting with a novel approach of PDE-based inpainting in q-space, which is the space spanned by diffusion sensitizing gradient directions and magnitudes. We find that predictions from linear diffusion in q-space can be made with low computational effort and are strong enough to further improve compression rates.

The remainder of our work is organized as follows: Sect. 2 provides the required background and discusses prior work on 4D image compression. Section 3 introduces the components of our proposed codec. Section 4 demonstrates that the resulting compression rates exceed those of several baselines and investigates the effects of specific design choices. Section 5 concludes with a brief discussion.

## 2 Background and Related Work

We will now introduce the main ideas behind diffusion PDE-based image inpainting and compression (Sect. 2.1), clarify the foundations of diffusion MRI and q-space (Sect. 2.2), and briefly review the literature on 4D medical image compression (Sect. 2.3).

### 2.1 Diffusion PDE-Based Inpainting and Compression

Inspired by their use for modeling physical phenomena, partial differential equations (PDEs) have a long tradition for solving problems in image processing. In particular, the PDE describing heat diffusion has provided a framework for image smoothing and inpainting [17–23].

The heat equation captures the relationship between temporal changes in a temperature  $\partial_t u$  and the divergence of its spatial gradient  $\nabla u$ ,

$$\partial_t u = \operatorname{div}(D \cdot \nabla u), \quad (1)$$

where  $D$  is the thermal diffusivity of the medium. In a homogeneous and isotropic medium, the diffusivity  $D$  is a constant scalar. In a non-homogeneous isotropic medium,  $D$  would still be a scalar, but depend on the spatial location. In an anisotropic medium, heat dissipates more rapidly in some directions than in others. In that case,  $D$  is a symmetric positive definite matrix that is referred to as a diffusion tensor.

In image processing, the gray value at a certain location is interpreted as a temperature  $u$ , and Eq. (1) is coupled with suitable boundary conditions. For image smoothing,

$$\begin{aligned} \partial_t u &= \operatorname{div}(D \cdot \nabla u), & \Omega \times (0, \infty), \\ \partial_n u &= 0, & \partial\Omega \times (0, \infty), \end{aligned} \quad (2)$$

where  $\Omega$  is the image domain and  $n$  is the normal vector to its boundary  $\partial\Omega$ . The original image  $f : \Omega \rightarrow \mathbb{R}$  is used to specify an initial condition  $u = f$  at  $t = 0$ . For increasing diffusion time  $t$ ,  $u$  will correspond to an increasingly smoothed version of the image.

In image inpainting, values are known at a subset of pixel locations, and unknown values should be filled in. For this, a Dirichlet boundary condition is introduced, which fixes values at a subset  $K$  of pixel locations [2,20]

$$\begin{aligned} \partial_t u &= \operatorname{div}(D \cdot \nabla u), & \Omega \setminus K \times (0, \infty), \\ \partial_n u &= 0, & \partial\Omega \times (0, \infty), \\ u &= f, & K \times [0, \infty). \end{aligned} \quad (3)$$

and a steady-state is computed at which  $\partial_t u \approx 0$ . The ability of PDEs to reconstruct plausible images even from a very sparse subset of pixels made them useful for image compression [2–4].

Different choices of diffusivity  $D$  introduce considerable flexibility with respect to shaping the final result. Fixing  $D = 1$  turns Eq. (3) into second-order linear homogenous (LH) diffusion

$$\partial_t u = Lu, \quad \Omega \setminus K \times (0, \infty) \quad (4)$$

with  $Lu = \Delta u$ , where  $\Delta$  denotes the Laplace operator, and the steady state satisfies the Laplace equation  $\Delta u = 0$ . Even though the resulting reconstructions suffer from singularities [2] and can often be improved by the more complex models discussed below, they have been used to design compression codecs for cartoon-like images [24], flow fields [25], and

depth maps [26–28]. Its simple linear nature and fast convergence to the steady-state also make LH diffusion an attractive choice for real-time video compression [5,6].

Compared to LH diffusion, decreasing the diffusivity as a function of image gradient magnitude permits a better preservation of salient edges [18,29]. This is referred to as nonlinear diffusion, since the results are no longer linear in the original image  $f$ . Rather than just decreasing the overall diffusivity close to edges, modeling  $D$  as an anisotropic diffusion tensor permits smoothing along edges, while maintaining or even increasing the contrast perpendicular to them. One widely used model is referred to as edge-enhancing diffusion (EED) [30].

All PDEs that have been discussed up to this point are of second order. Fourth- and higher-order extensions have also been studied, both for smoothing [31–35] and for inpainting [36,37]. In the simplest case, setting  $Lu = -\Delta^2 u$  in Eq. (4) leads to the biharmonic (BH) equation. In two and three dimensions, it does not suffer from the singularities that are present in the results of LH diffusion [2,38], while preserving a simple linear nature. For this reason, BH has been considered for the design of compression codecs [38–41]. However, it no longer satisfies a min-max principle [33] and it increases running time and sensitivity to quantization error.

Our own previous work [37] proposed an anisotropic fourth-order PDE in which a fourth-order diffusion tensor is constructed from the image gradient in a similar way as in second-order EED. We thus refer to it as fourth-order edge-enhancing diffusion (FOEED). It was shown to result in more accurate inpainting results than second-order EED, and higher PDE-based compression rates, in several examples [12].

Our current work is concerned with compressing data from diffusion MRI, which is similar to hyperspectral images in that it contains a large number of values (channels) at each location [40]. However, the channels in hyperspectral images have a one-dimensional order, while our channels correspond to positions on a spherical shell in  $q$ -space, which we will now introduce.

## 2.2 Diffusion MRI

The signal in magnetic resonance imaging is generated by the hydrogen atoms within water molecules. Their heat motion is referred to as self-diffusion, since it takes place despite a zero concentration gradient. The extent to which this motion is restricted in a cellular environment correlates with microstructural parameters such as cellular density or integrity. Moreover, in the white matter of the human brain, which contains the tracts that connect different brain regions, self-diffusion can occur more freely in the local direction of those tracts than perpendicular to it [42]. Therefore, measuring the apparent diffusion coefficient in different directions

provides relevant information about small-scale structures that are below the image resolution of in vivo MRI.

This motivates the use of diffusion MRI. It goes back to the idea of measuring diffusion by introducing a pair of diffusion sensitizing magnetic field gradients into a nuclear magnetic resonance sequence [43]. Integrating it with spatially resolved magnetic resonance imaging permits diffusion measurements at a voxel level [13]. Repeating the measurements with differently oriented gradients reveals a biologically relevant directional dependence in various tissue types, including muscle and the white matter of the brain [44].

Several key parameters of the diffusion sensitization can be summarized in the gradient wave vector

$$\mathbf{q} = \frac{1}{2\pi} \gamma \delta \mathbf{g}, \quad (5)$$


where  $\gamma$  is the gyromagnetic ratio of hydrogen nuclei in water,  $\delta$  is the duration of the diffusion sensitizing gradients, and  $\mathbf{g}$  corresponds to their direction and strength. The normalized MR echo magnitude  $|E(\mathbf{q}, \tau)|$  additionally depends on the time  $\tau$  between the pair of gradient pulses. It is computed as the ratio between the corresponding diffusion-weighted measurement and an unweighted measurement with  $\mathbf{q} = \mathbf{0}$ . It is antipodally symmetric,  $|E(-\mathbf{q}, \tau)| = |E(\mathbf{q}, \tau)|$ .

The relevance of this  $q$ -space formalism derives from a Fourier relationship between  $|E(\mathbf{q}, \tau)|$  and the ensemble average diffusion propagator  $\bar{P}(\mathbf{R}, \tau)$ , which specifies the probability of a molecular displacement  $\mathbf{R}$  within a fixed diffusion time [45]. An alternative parameterization of the diffusion gradients is in terms of their direction and a factor  $b = 4\pi^2 \|\mathbf{q}\|^2 (\tau - \delta/3)$ , which also accounts for the fact that the diffusion weighting increases with the effective diffusion time  $(\tau - \delta/3)$ .

Due to practical constraints on the overall duration of dMRI measurements, the sampling of  $q$ -space is usually limited to one or several reference measurements with  $\mathbf{q} = \mathbf{0}$ , as well as one or a few shells with constant  $\|\mathbf{q}\|$ , and thus constant  $b$ . This is illustrated in Fig. 1. Such setups focus on the directional dependence of the signal, and typically strive for a uniform distribution of gradient directions on these shells [46]. Our codec assumes dMRI data with such a “shelled” structure, an assumption that is shared by well-established algorithms in the field [47].

## 2.3 4D Medical Image Compression

Many medical imaging modalities, including magnetic resonance imaging, computed tomography, and ultrasound, can be used to image volumes repeatedly, in order to capture time-dependent phenomena such as organ motion, perfusion, or blood oxygenation. Considerable work has been done on lossless and lossy compression of the resulting 4D (3D plus



**Fig. 1** Illustration of three diffusion sensitizing gradient directions on a shell in  $q$  space, with equal  $b = 700$  (top left). The three diffusion-weighted images have been measured with different gradient directions, as indicated at the bottom left of each image. Comparing them reveals the directional dependence of the dMRI signal

time) image data. Much of it has borrowed from video coding and has often involved motion compensation [48,49], which is combined with wavelet transforms [48,50–53] or hierarchical vector quantization [54] for compression.

Almost all of these works have compared their compression rates to codecs from the JPEG family. We will also compare our codec to JPEG-LS, lossless JPEG, and JPEG2000. Additionally, we compare compression rates against GZIP [16] which, in conjunction with the Neuroimaging Informatics Technology Initiative (NIfTI) file format, is currently most widely used to compress diffusion MRI data in practice. To make this comparison fair, we also use Huffman coding or Deflate within our own codec, as opposed to computationally efficient alternatives that might further improve compression rates [55,56].

Even though the volumetric images in diffusion MRI are also taken sequentially, their temporal order is less relevant than the  $q$ -space structure that was described above: Measuring with the same diffusion sensitizing gradients, but in a different order, should yield equivalent information, even though it permutes the temporal order. To the best of our knowledge, no codec has been proposed so far that exploits this very specific structure. There has been extensive work on compressed sensing for diffusion MRI (see [57,58] and references therein), but with a focus on reducing measurement time, rather than efficient storage of the measured data.

Recent work has demonstrated the potential of deep learning for lossless compression of 3D medical images [59].

Extending this specifically for diffusion MRI is an interesting future direction. However, our PDE-based approach has the advantage of not requiring any training data. Since medical data are a particularly sensitive type of personal data, obtaining diverse large-scale datasets can be difficult, and the potential of model attacks that could cause data leakage is concerning [60,61].

### 3 Proposed Lossless Codec

Traditional PDE-based image compression [2,3,12,37] performs inpainting in image space, which relies on piecewise smoothness of the image. A key contribution of our current work is to additionally exploit the smoothness in  $q$ -space. As it can be seen in Fig. 1, dMRI signals that are measured with similar gradient directions are correlated.

Our codec uses a spatial PDE for the first few volumes, which is described in more detail in Sect. 3.1. Once sufficiently many samples are available so that a  $q$ -space PDE, described in Sect. 3.2, produces stronger compression than the spatial PDE, we switch to it.

The  $q$ -space PDE assumes that all volumes are in correct spatial alignment, which might be violated in practice due to subject motion. For this reason, our codec includes a mechanism for motion compensation, described in Sect. 3.3. Our overall compressed file format is specified in Sect. 3.4.

#### 3.1 Lossless 3D Spatial Codec

The initial few volumes are compressed with an image space PDE-based codec that follows our recent conference paper [12]. To make our current work self-contained, we briefly summarize the most relevant points, focusing on the forward, i.e., encoding direction. The decoding process just mirrors the respective steps. The codec is composed of three main parts: data sparsification (initial mask selection), prediction (iterative reconstruction), and residual coding.

*Initial Mask Selection* As an initial mask, our codec simply stores voxel intensities on a sparse regular grid. More specifically, for a given 3D input image of size  $n_1 \times n_2 \times n_3$ , the initial mask is chosen as a hexahedral grid consisting of voxels  $(4i_1, 4i_2, 4i_3)$ , where  $i_j \in \{0, 1, \dots, \lfloor \frac{n_j-1}{4} \rfloor\}$ ,  $j \in 1, 2, 3$ .

Most lossy PDE-based codecs select a mask adaptively [2,3], which better preserves important image features such as edges and corners [24]. However, this introduces the need to store the locations of the selected pixels, which can be avoided by the use of fixed grids [27,62]. In the context of lossless compression, we achieved higher compression rates by combining the latter strategy with iterative reconstruction.

*Iterative Reconstruction* Making PDE-based compression lossless requires coding the differences between the original image and the PDE-based reconstruction and is beneficial in terms of compression rates to the extent that those residuals are more compressible than the original image. In general, residuals become more compressible the more accurate the reconstruction is. Therefore, the overall compression rate can be increased by iteratively coding residuals of some pixels, and refining the remaining ones based on them.

Our previous work [12] explored different iterative schemes. The variant that is used here codes the residuals in all remaining face-connected neighbors of the current mask voxels, i.e., up to six voxels per mask voxel. Those neighbors become part of the mask for the next iteration, and the process continues until all voxels have been coded.

Among the PDEs that have been explored for inpainting, we currently consider the two that worked best in [12], i.e., traditional edge-enhancing diffusion (EED) [20] and our recent fourth-order generalization (FOEED) [37].

*Residual Coding* Residuals are computed in modular arithmetic, so that they can be represented as unsigned integers. The final compression of the initial mask and the residuals is either done via a Huffman entropy encoder or the Deflate algorithm, depending on which gives the smaller output file size.

In cases where medical images contain a substantial amount of empty space, e.g., a background region with exactly zero image intensity, our previous work [12] found that coding it separately using run length encoding (RLE) can provide an additional benefit. Unfortunately, in dMRI, the background is perturbed by measurement noise, which renders this approach ineffective. Therefore, our current work does not include any dedicated empty space coding.

### 3.2 PDE-Based q-Space Inpainting

The general idea of q-space inpainting is illustrated in Fig. 2: Once a certain number of diffusion-weighted images with different gradient directions are known, we can use them to predict images that correspond to a new direction. This happens at the voxel level, so that the prediction at a given location is entirely determined by values at the same location in the known images.

This can be understood as “flipping” the setup from Sect. 3.1, where the mask consisted of pixel locations, and the inpainting was repeated with an identical mask for each channel. Instead, the mask now specifies the known channels, and inpainting is repeated for each voxel in the volume.

#### 3.2.1 Compressing Diffusion-Weighted Images

Since we assume that diffusion-weighted measurements are on spherical shells in q-space (Sect. 2.2), we inpaint with




Fig. 2 An example of q-space inpainting to predict a diffusion-weighted volume in gradient direction  $\mathbf{g}_5$  (red double arrow) based on four known volumes, illustrated as filled volumes in directions  $\mathbf{g}_i$  (black arrows) (Color figure online)




Fig. 3 The q space sampling of the dMRI data used in our experiments (left), and the resulting triangulation that is used within the finite element method

second-order linear homogeneous (LH) diffusion

$$\partial_t u = \Delta u \tag{6}$$

or fourth-order biharmonic (BH) smoothing

$$\partial_t u = -\Delta^2 u \tag{7}$$

on the sphere, where  $\Delta$  is the Laplace-Beltrami operator.

Given that our samples do not form a regular grid, we numerically solve these equations using finite element methods (FEMs) [63,64]. For this, we first construct a 3D Delaunay tessellation from the set of all gradient vectors  $\mathbf{g}_i$

and their antipodal counterparts  $-\mathbf{g}_i$ , and then extract a triangular surface mesh from it. Figure 3 shows an example of the given vectors (left), and the resulting triangular mesh (right).

Similar to PDE-based inpainting in the image domain, we fix the known values by imposing Dirichlet boundary conditions at the vertices corresponding to the previously coded diffusion-weighted images, again accounting for antipodal symmetry so that each known image determines the values of two vertices. Once a steady state has been reached, the values at locations corresponding to diffusion-weighted images that are yet to be coded can serve as predictions. Similar as before, we compute residuals with respect to those predictions in modular arithmetic and apply Huffman coding or Deflate to them.

We found that once a sufficient number of diffusion-weighted images are available as a basis of q-space inpainting, its residuals become more compressible than those from iterative image space inpainting. Our codec adaptively determines a suitable point for switching from spatial to q-space predictions. After the first diffusion-weighted volume, it starts comparing the sizes of compressing subsequent volumes with the spatial codec (Sect. 3.1) to the size when using q-space inpainting and switches to it on the first volume where it is beneficial. To limit computational effort, the spatial codec is no longer tried for subsequent volumes.

### 3.2.2 Accelerated Computation

Since q-space inpainting happens at a voxel level, it should be repeated for each voxel of the 3D image. However, the computational cost of running the FEM solver for each voxel separately is extremely high. Fortunately, linearity of the PDEs and the fact that the Dirichlet boundary conditions are imposed on the same vertices for each voxel permit a significant speedup.

Formally, we can consider one time step of numerically solving Eqs. (6) or (7), at time  $t$ , as applying a discrete linear differential operator  $\mathbf{D}$ , which is determined by the vertices and connectivity of our triangular mesh, on a discrete input  $\mathbf{u}^{(t)} \in \mathbb{R}^c$ ,

$$\mathbf{u}^{(t+1)} = \mathbf{D} \left[ \mathbf{u}^{(t)} \right], \tag{8}$$

where  $c$  is the number of channels (q-space samples per voxel). The boundary conditions ensure that  $u_{k_j}^{(t+1)} = u_{k_j}^{(t)}$  at positions  $k_j$  that correspond to the fixed (previously coded) channels.

The inpainting result is obtained as the fixed point  $\mathbf{u}^{(\text{FP})}$  as  $t \rightarrow \infty$ . It can be approximated by iterating  $\mathbf{D}$  a sufficient number of times, resulting in an operator  $\mathbf{D}_{\text{FP}}$  that directly

maps

$$\mathbf{u}^{(\text{FP})} = \mathbf{D}_{\text{FP}} \left[ \mathbf{u}^{(0)} \right]. \tag{9}$$

$\mathbf{D}_{\text{FP}}$  is still linear, and we observe that its kernel is the subspace corresponding to the unknown q-space samples, so that their initialization in  $\mathbf{u}^{(0)}$  does not influence the steady state [4]. Therefore, we can rewrite Eq. 9 as

$$\begin{aligned} \mathbf{D}_{\text{FP}} \left[ \mathbf{u}^{(0)} \right] &= u_{k_1}^{(0)} \mathbf{D}_{\text{FP}}[\mathbf{e}_{k_1}] + u_{k_2}^{(0)} \mathbf{D}_{\text{FP}}[\mathbf{e}_{k_2}] \\ &+ \dots + u_{k_n}^{(0)} \mathbf{D}_{\text{FP}}[\mathbf{e}_{k_n}], \end{aligned} \tag{10}$$

where  $\mathbf{e}_{k_j}$  are the indicator vectors of the known samples  $k_j$ .

In other words, by computing  $\mathbf{w}_{k_j} = \mathbf{D}_{\text{FP}}[\mathbf{e}_{k_j}]$ , we can obtain weight vectors that specify how the known values are combined to predict the unknown ones. They are analogous to the ‘‘inpainting echoes’’ that have been computed in previous work [65] for the purpose of optimizing tonal data. Omitting the irrelevant initialization of the unknown values from the input  $\mathbf{u}^{(0)}$ , and the known values from the output  $\mathbf{u}^{(\text{FP})}$  yields a weight matrix  $\mathbf{W}$  of shape  $m \times n$  for  $n$  known and  $m$  unknown values.

We compute the coefficients of  $\mathbf{W}$  by running the FEM  $n$  times. In the  $j$ th run, we set the value corresponding to  $k_j$  to one, all remaining values to zero. After numerically solving the Laplace or Biharmonic PDE, the values at the  $m$  unique vertices that correspond to unknown DWIs yield the  $j$ th column of  $\mathbf{W}$ .

Finally,  $\mathbf{W}$  allows us to make efficient predictions in each voxel, by simply multiplying it to a vector that contains the intensities in that voxel from the previously coded diffusion gradients.

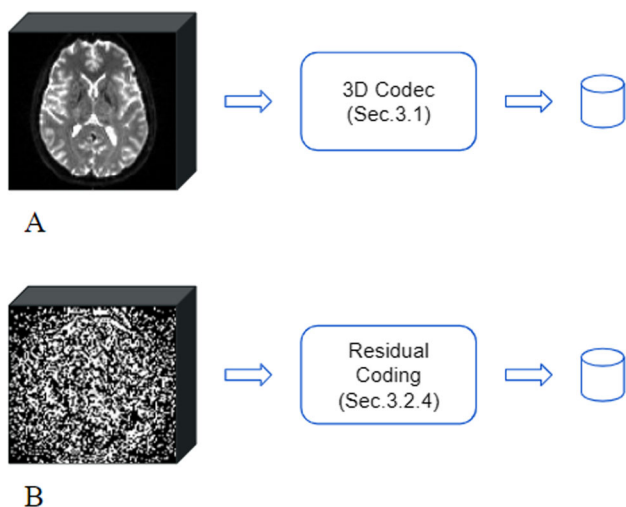
### 3.2.3 Implementation Details and Running Times

We numerically solve Eqs. (6) and (7) via the open-source FEM solver package FEniCSx [64]. For implementation details, we refer to its tutorials [66]. Applying this solver to each voxel of a  $104 \times 104 \times 72$  volume takes close to two and four hours, respectively, for LH and BH PDEs on a single 3.3 GHz CPU core.

The acceleration from the previous section reduces this to only 1.6 s and 2.4 s per volume, respectively. This includes the time for building a Delaunay tessellation, which is computed with SciPy [67], and extracting a surface mesh from it using the BoundaryMesh method from FEniCSx.

### 3.2.4 Compressing $b = 0$ Images

Our general approach simplifies for unweighted volumes with  $b = 0$ . Again, the first of them is compressed using the spatial codec. If multiple  $b = 0$  images were acquired



**Fig. 4** Top **A**: Compression of the first  $b = 0$  volume using the recently proposed lossless 3D codec [12]. Bottom **B**: Compression of the remaining  $b = 0$  volumes using residuals in modular arithmetic. Residuals are taken with respect to the first  $b = 0$  volume, after motion correction


to increase the signal-to-noise ratio, our codec compresses the remaining ones by taking the respective residuals with respect to the first  $b = 0$  volume, as illustrated in Fig. 4.

### 3.3 Motion Correction

Subject motion commonly occurs during the lengthy dMRI acquisitions and is typically accounted for by applying motion correction based on image registration [68]. We also include this step in our codec, since inpainting in q-space requires a correct spatial alignment of all 3D volumes so that predictions are based on information from the same location within the brain.

We implement motion correction as follows:

1. We perform affine registration of each volume to the same  $b = 0$  volume, which is used as a common reference. This yields a transformation matrix  $\mathbf{T}_{X \rightarrow b_0}$  which aligns DWI volume  $X$  to the  $b = 0$  reference.
2. When predicting a DWI volume  $P$ , we transform all known volumes  $X$  via the affine transformation  $\mathbf{T}_{P \rightarrow b_0}^{-1} \circ \mathbf{T}_{X \rightarrow b_0}$ , which can be computed from the transformations in Step 1.
3. In addition to resampling each known volume  $X$ , we reorient its diffusion gradient direction  $\mathbf{g}_X$  according to the rotational part of the transformation in Step 2. Omitting this step would lead to incorrect relative orientations of diffusion gradient directions [69], which could again reduce accuracy of q-space inpainting.



**Fig. 5** Boundary effects in volume alignment. Top: Original DWI volume. Center and Bottom: Motion corrected with zero padding and nearest neighbor extrapolation, respectively

Transforming images via a common reference allows us to align them without having to perform image registration on all pairs of volumes. This saves considerable computational effort. Combining the two transformations and applying them in a single step also reduces computational effort, and simultaneously reduces image blurring compared to a two-stage implementation that would involve interpolating twice.

In addition to the computational expense, motion correction incurs the cost of having to store the affine matrices  $\mathbf{T}_{X \rightarrow b_0}$  along with the compressed data. Experimental results in Sect. 4.4 will demonstrate that this storage cost is outweighed by the increase in compression rate when q-space inpainting properly accounts for motion.

Subject movement correction and B-matrix reorientation are done using the freely available FSL tools [68] and the DIPY imaging library [70], respectively. A practically relevant implementation detail concerns boundary effects. As illustrated in Fig. 5, missing information can enter the field of view when applying image transformations. We found that q-space inpainting near the boundary of the domain works more reliably if we resolve these cases with nearest neighbor extrapolation, rather than with zero padding.

### 3.4 Compressed File Format

In our current implementation, the relevant data are spread over multiple files whose sizes are added when computing compression rates.

The volumes that are compressed with the 3D lossless codec (Sect. 3.1) are stored with the same header as in [12]. Stated briefly, it contains the original minimum and maximum voxel values (4 bytes), sizes of the compressed data streams for zero voxel binary mask and mask intensities (8 bytes), the diffusivity contrast parameter (4 bytes), the type of PDE (2 bits), the dilation mode (1 bit), and the types of encoding for mask intensities and residuals (2 bits).

For each volume that is compressed with q-space inpainting, the header contains the original minimum and maximum voxel values (4 bytes), the type of PDE (2 bits), the type of encoding for the residuals (1 bit), and the volume number in the original order (2 bytes).

Mask and residual values themselves are stored after compression with pure Huffman coding or Deflate, depending on what gave a smaller file size.

In addition, we store the NIFTI header (348 bytes) as well as files containing  $b$  values and gradient vectors in their original ASCII formats. For simplicity, affine transformations for motion correction are also kept in the ASCII format generated by FSL FLIRT [68].

## 4 Results and Discussion

### 4.1 Data

We evaluate our codec on two dMRI datasets that were made publicly available by Koch et al. [71] and are specifically suited to investigate the impact of subject motion compensation. Both datasets have been collected from the same subject (male, 34 years) in the same scanner, a 3T MAGNETOM Skyra (Siemens Healthcare, Erlangen, Germany), with an identical measurement sequence. For the first scan, the subject received the usual instruction of staying as still as possible during the acquisition. For the second scan, the subject was asked to move his head, to deliberately introduce motion artifacts.

From these datasets, we use the five non-diffusion weighted ( $b = 0$ ) MRI scans each, as well as 30 diffusion weighted images ( $b = 700 \text{ s/mm}^2$ , diffusion gradient duration  $\delta = 334 \text{ ms}$ , spacing  $\tau = 445 \text{ ms}$ ). Each image consists of  $104 \times 104 \times 72$  voxels with a resolution of  $2 \times 2 \times 2 \text{ mm}^3$ . The data, and the effects of subject motion, are illustrated in Fig. 6.

### 4.2 DTI Baseline

We compare the signal predictions from our q-space PDE to a simple baseline, which is derived from the diffusion tensor imaging (DTI) model. DTI is widely used in practice, due to its relative simplicity and modesty in terms of scanner hardware and measurement time.

It rests on the assumption that the diffusion propagator  $\bar{P}(\mathbf{R}, \tau)$  is a zero-mean Gaussian whose covariance matrix is proportional to the diffusion tensor  $\mathbf{D}$ , a symmetric positive definite  $3 \times 3$  matrix that characterizes the local diffusion [14]. The signal model in DTI relates the diffusion-weighted signal  $S(\hat{\mathbf{g}}, b)$  for a given  $b$ -value and gradient vector direction  $\hat{\mathbf{g}} = \mathbf{g}/\|\mathbf{g}\|$  to the unweighted signal  $S_0$  according to

$$S(\hat{\mathbf{g}}, b) = S_0 e^{-b \hat{\mathbf{g}}^T \mathbf{D} \hat{\mathbf{g}}}. \quad (11)$$

Fitting  $\mathbf{D}$  requires at least one reference MR image  $S_0$ , plus diffusion-weighted images in at least six different directions, which are usually taken with a fixed nonzero  $b$ -value. Equation (11) can then be used to predict the diffusion-weighted signal in any desired direction. In our experiments, we compare our PDE-based to DTI-based predictions that account for the same set of known measurements.

### 4.3 Comparing Lossless Codecs for Diffusion MRI


A comparison of file sizes that can be achieved on our two test datasets with different lossless codecs is provided in Table 1. As a baseline, the first two rows show results from coding each 3D volume independently with our recently proposed PDE-based codec [12], using second-order (R-IEED-1) and fourth-order anisotropic diffusion (R-IFOEED-1). Additional savings of other codecs with respect to R-IEED-1 are given in percent.

The second block in Table 1 shows results from several variants of our proposed new codec, which adaptively combines inpainting in q-space and image space. Highest compression rates were achieved when combining linear homogeneous (LH) diffusion in q-space with R-IFOEED-1 in image space, closely followed by R-IEED-1. Biharmonic (BH) diffusion in q-space also produced useful, but slightly weaker results.

Both q-space diffusion approaches achieved better compression than predictions from DTI (Sect. 4.2). This could be due to the fact that the quadratic model of diffusivities in Equation (11) is known to be an oversimplification in many parts of the brain [72], and the PDE-based approaches provide more flexibility.

DTI requires independent coding of at least seven 3D volumes, which led us to fix this split in our experiments. PDE-based imputation makes it possible to switch to q-space inpainting earlier, and our adaptive selection does so after





**Fig. 6** Example images from our two dMRI datasets, without deliberate head motion (left) and with strong motion artifacts (right). In each case, six corresponding sagittal slices from different diffusion weighted

images (DWIs) are shown. Note that subject motion leads to spatial misalignments between DWIs, but also to artifacts within individual images

**Table 1** Compressed file sizes from separate PDE-based compression of each 3D scan (baseline), from different variants of our proposed lossless codec, as well as from GZIP and lossless codecs from the JPEG family

Codec variant	Scan 1: No deliberate motion			Scan 2: Strong head motion		
	Split	Size (bytes)	Over R-IEED-1 (%)	Split	Size (bytes)	Over R-IEED-1 (%)
R-IEED-1		16022666	–		16082537	–
R-IFOEED-1		15955826	+ 0.42		16019913	+ 0.39
qLH ◦ R-IFOEED-1	27/4	14984472	+ 6.50	26/5	15570493	+ 3.18
qLH ◦ R-IEED-1	27/4	14991732	+ 6.43	26/5	15578604	+ 3.13
qBH ◦ R-IFOEED-1	27/4	15032384	+ 6.18	26/5	15681354	+ 2.50
qBH ◦ R-IEED-1	27/4	15039644	+ 6.14	26/5	15689465	+ 2.44
DTI ◦ R-IFOEED-1	24/7	15099213	+ 5.76	24/7	15854216	+ 1.42
DTI ◦ R-IEED-1	24/7	15108244	+ 5.71	24/7	15861176	+ 1.38
GZIP		21841701	– 36.32		21819641	– 35.67
JPEG		17885953	– 11.63		17905933	– 11.34
JPEG-LS		17921807	– 11.85		17893931	– 11.26
JPEG2000		15993453	+ 0.18		15980005	+ 0.64

For hybrid codecs, the split indicates the number of volumes coded with q-space or spatial inpainting, respectively

four volumes in the low-motion data, after five volumes in the data with strong motion.

Switching to q-space inpainting also speeds up our codec. Our implementation of R-IEED-1 and R-IFOEED-1 requires approximately 478 s and 6185 s, respectively, for one volume on a single 3.3 GHz CPU core. Even though it would be possible to further optimize this, exploiting linearity in qLH and qBH, as described in Sect. 3.2, significantly lowers the intrinsic computational complexity, so that even a straightforward implementation only requires 1.64 s and 2.4 s per volume, respectively.

It can be seen in Fig. 6 that subject motion during different phases of the acquisition leads to different types of artifacts. Results in Table 1 include the motion correction described in Sect. 3.3, which compensates spatial misalignments of different scans. However, motion can also lead to signal dropouts or to distortions within scans, which our current codec does not explicitly account for. This explains why q-space inpainting is less effective on the second as compared to the first

scan. However, even on this challenging dataset that exhibits unusually strong artifacts, q-space inpainting still provides a benefit compared to all other alternatives.


Finally, Table 1 shows results from several other lossless codecs for comparison. GZIP is most widely used in practice, but the resulting files are more than 35% larger than those from our proposed codec. Among the lossless codecs from the JPEG family, JPEG2000 is the only one that outperforms R-IEED-1 for per-volume compression, and only by a small margin. Our new hybrid methods that combine image space and q-space inpainting always performed best.

#### 4.4 Benefit from Motion Correction

Table 2 investigates the benefit of motion correction (Sect. 3.3) by showing file sizes when removing motion correction from our codec, and comparing the results to ones with motion correction (Table 1), indicating the benefit in percent.

**Table 2** Compressed file sizes when omitting motion compensation, and the relative benefit from motion correction

Codec variant	Scan 1: No deliberate motion			Scan 2: Strong head motion		
	Split	Without correction	Benefit (%)	Split	Without correction	Benefit (%)
qLH ◦ R-IFOEED-1	27/4	15113407	+ 0.85	16/15	16456557	+ 5.38
qLH ◦ R-IEED-1	27/4	15122438	+ 0.86	16/15	16478088	+ 5.46
qBH ◦ R-IFOEED-1	27/4	15302805	+ 1.77	16/15	16582150	+ 5.43
qBH ◦ R-IEED-1	27/4	15311836	+ 1.78	16/15	16603681	+ 5.51
DTI ◦ R-IFOEED-1	24/7	15396194	+ 1.93	24/7	16946648	+ 6.45
DTI ◦ R-IEED-1	24/7	15405225	+ 1.93	24/7	16955494	+ 6.45



**Fig. 7** Given a set of previously coded DWIs, the *closest* strategy (left) selects the volume whose gradient vector has the smallest angular distance from the known ones, to maximize expected prediction accuracy. The *furthest* strategy (right) maximizes the angular distance, aiming for a more uniform coverage of the sphere for subsequent steps. The sketch shows the directions selected in the first three steps as black double arrows, the fourth direction as a red dot

Even on the first scan, in which the subject tried to keep his head still, compensating for small involuntary movements yields a slight benefit. The effect is largest when imputing via qBH and DTI. This might be explained by the fact that qLH satisfies the min-max principle, which makes it more robust against inaccuracies in its inputs, and provides another argument in its favor.

When strong head motion is present (second scan), restoring a correct voxel alignment via motion correction becomes essential for q-space inpainting. Without it, the switch to q-space imputation happens much later, and the overall file size is larger than when coding each volume independently. This is explained by the fact that our codec always applies difference coding to the  $b = 0$  images and that this becomes detrimental when those images are strongly misaligned.

### 4.5 Effect of Re-ordering DWIs

Since q-space imputation relies on the previously (de)coded diffusion weighted images, its accuracy depends on the order in which we process the gradient directions.

Two contradictory greedy strategies are illustrated in Fig. 7: Always selecting the *closest* gradient direction, i.e.,

**Table 3** Compressed file size for scan 1 (without strong motion) when ordering the diffusion-weighted images differently

Codec variant	Closest selection	Furthest selection
qLH ◦ R-IEED-1	15031703	14991732
qBH ◦ R-IEED-1	15153333	15039644
qDTI ◦ R-IEED-1	15290105	15108244

The ordering affects the accuracy of q-space imputation

the one with the smallest angular distance from the already known ones, can be expected to result in the most accurate prediction, in the same spirit as our spatial codec (Sect. 3.1) iteratively grows a mask of known pixels around an initial set of seed points.

On the other hand, the spatial codec starts with a seed mask that covers the full domain sparsely, but uniformly. Achieving something similar motivates selecting the gradient direction that is *furthest* from any of the known ones. Even though this strategy can be expected to lead to lower accuracy, and therefore to less compressible residuals in the first few iterations, later iterations might benefit from the more uniform coverage of the overall (spherical) domain.

Table 3 presents the effect of these two selection strategies on final file sizes. The results are from the first scan, without strong motion. Overall, greedily selecting the furthest gradient vector gives slightly smaller overall file sizes. Therefore, this is the strategy that we followed in all other experiments.

## 5 Conclusion

In this work, we introduced a PDE-based lossless image compression codec that explicitly exploits both the spatial and the q-space structure in diffusion MRI. To our knowledge, it is the first codec that has been tailored to this type of data. We demonstrated a clear improvement over PDE-based codecs that treat each volume separately, and over other established baselines including GZIP and spatial codecs from the JPEG family.

We evaluated several variants of our codec, and found that q-space predictions with linear homogeneous diffusion permitted the highest compression rates among them. With our proposed method for accelerated computation, it could also be applied at a very reasonable computational cost. We further demonstrated the importance of including motion correction, and propose an efficient implementation that is based on affine image transformations via a common reference. Finally, we found that the order of coding the diffusion-weighted volumes had a relatively minor effect, but that a greedy strategy that strives to cover the sphere as uniformly as possible provides a small benefit.

In the future, one might attempt to replace the switching between image space and q-space inpainting with a PDE that jointly operates on the product space. However, this is likely to substantially increase the computational effort, and introduces the issue of properly balancing image space and q-space diffusion. Similarly, employing nonlinear PDEs for q-space predictions might further increase compression rates, but is likely to cause a high computational cost.

**Acknowledgements** This work has been supported by the German Academic Exchange Service (DAAD).

**Funding** Open Access funding enabled and organized by Projekt DEAL.

## Declarations

**Conflict of interest** The authors declare that they have no conflict of interest.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Dinov, I.D.: Volume and value of big healthcare data. *J. Med. Stat. Inf.* **4**, 3 (2016)
- Galić, I., Weickert, J., Welk, M., Bruhn, A., Belyaev, A., Seidel, H.-P.: Image compression with anisotropic diffusion. *J. Math. Imaging Vis.* **31**(2–3), 255–269 (2008)
- Schmaltz, C., Peter, P., Mainberger, M., Ebel, F., Weickert, J., Bruhn, A.: Understanding, optimising, and extending data compression with anisotropic diffusion. *Int. J. Comput. Vis.* **108**(3), 222–240 (2014)
- Peter, P., Schmaltz, C., Mach, N., Mainberger, M., Weickert, J.: Beyond pure quality: Progressive modes, region of interest coding, and real time video decoding for PDE-based image compression. *J. Vis. Commun. Image Represent.* **31**, 253–265 (2015)
- Köstler, H., Stürmer, M., Freundl, C., Rüdte, U.: PDE based video compression in real time. Technical Report 07–11, University Erlangen–Nürnberg, Lehrstuhl für Informatik 10 (2007)
- Andris, S., Peter, P., Weickert, J.: A proof-of-concept framework for PDE-based video compression. In: *Proc. Picture Coding Symposium (PCS)*, pp. 1–5 (2016). IEEE
- Pennebaker, W.B., Mitchell, J.L.: *JPEG Still Image Data Compression Standard*, 1st edn. Kluwer Academic Publishers, USA (1992)
- Taubman, D., Marcellin, M.: *JPEG2000: Image Compression Fundamentals, Standards and Practice*, 1st edn. Springer, USA (2002)
- Peter, P.: Three-dimensional data compression with anisotropic diffusion. In: *German Conf. on Pattern Recognition*, pp. 231–236 (2013). Springer
- Kil, S.-K., Lee, J.-S., Shen, D., Ryu, J., Lee, E., Min, H., Hong, S.: Lossless medical image compression using redundancy analysis. *Int. J. Comput. Sci. Netw. Secur.* **6**(1), 50–56 (2006)
- Miaou, S.-G., Ke, F.-S., Chen, S.-C.: A lossless compression method for medical image sequences using JPEG-LS and inter-frame coding. *IEEE Trans. Inf. Technol. Biomed.* **13**(5), 818–821 (2009)
- Jumakulyyev, I., Schultz, T.: Lossless PDE-based compression of 3D medical images. In: *International Conference on Scale Space and Variational Methods in Computer Vision*, pp. 450–462 (2021). Springer
- Le Bihan, D., Breton, E., Lallemand, D., Grenier, P., Cabanis, E., Laval-Jeantet, M.: MR imaging of intravoxel incoherent motions: Application to diffusion and perfusion in neurologic disorders. *Radiology* **161**(2), 401–407 (1986)
- Basser, P.J., Mattiello, J., Le Bihan, D.: Estimation of the effective self-diffusion tensor from the NMR spin echo. *J. Magn. Reson.* **103**, 247–254 (1994)
- Jones, D.K. (ed.): *Diffusion MRI: Theory, Method, and Applications*, 1st edn. Oxford University Press, United Kingdom (2011)
- Deutsch, P.: *RFC1952: GZIP File Format Specification Version 4.3*. RFC Editor, USA (1996)
- Iijima, T.: Basic theory on the normalization of pattern (in case of typical one-dimensional pattern). *Bull. Electrotech. Lab.* **26**, 368–388 (1962)
- Perona, P., Malik, J.: Scale-space and edge detection using anisotropic diffusion. *IEEE Trans Pattern Anal. Mach. Intell.* **12**(7), 629–639 (1990)
- Osher, S., Rudin, L.I.: Feature-oriented image enhancement using shock filters. *SIAM J. Numer. Anal.* **27**(4), 919–940 (1990)
- Weickert, J.: *Anisotropic Diffusion in Image Processing*. Teubner, Stuttgart (1998)
- Masnou, S., Morel, J.-M.: Level lines based disocclusion. In: *Proc. Int'l Conf. on Image Processing (ICIP)*, pp. 259–263 (1998). IEEE
- Bertalmio, M., Sapiro, G., Caselles, V., Ballester, C.: Image inpainting. In: *Proc. Conf. on Computer Graphics and Interactive Techniques (SIGGRAPH)*, pp. 417–424 (2000)
- Chan, T.F., Shen, J.: Nontexture inpainting by curvature-driven diffusions. *J. Vis. Commun. Image Represent.* **12**(4), 436–449 (2001)
- Mainberger, M., Bruhn, A., Weickert, J., Forchhammer, S.: Edge-based compression of cartoon-like images with homogeneous diffusion. *Pattern Recogn.* **44**(9), 1859–1873 (2011)
- Jost, F., Peter, P., Weickert, J.: Compressing flow fields with edge-aware homogeneous diffusion inpainting. In: *IEEE Int'l Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 2198–2202 (2020)
- Gautier, J., Le Meur, O., Guillemot, C.: Efficient depth map compression based on lossless edge coding and diffusion. In: *2012 Picture Coding Symposium*, pp. 81–84 (2012). IEEE

27. Hoffmann, S., Mainberger, M., Weickert, J., Puhl, M.: Compression of depth maps with segment-based homogeneous diffusion. In: *Int'l Conf. on Scale Space and Variational Methods in Computer Vision*, pp. 319–330 (2013). Springer
28. Li, Y., Sjöström, M., Jennehag, U., Olsson, R.: A scalable coding approach for high quality depth image compression. In: *Proc. 3DTV-Conference: The True Vision-Capture, Transmission and Display of 3D Video (3DTV-CON)*, pp. 1–4 (2012). IEEE
29. Charbonnier, P., Blanc-Féraud, L., Aubert, G., Barlaud, M.: Deterministic edge-preserving regularization in computed imaging. *IEEE Trans. Image Process.* **6**(2), 298–311 (1997)
30. Weickert, J.: Theoretical foundations of anisotropic diffusion in image processing. In: Kropatsch, W., Klette, R., Solina, F., Albrecht, R. (eds.) *Theoretical Foundations of Computer Vision*, pp. 221–236. Springer, Vienna (1996)
31. You, Y.-L., Kaveh, M.: Fourth-order partial differential equations for noise removal. *IEEE Trans. Image Process.* **9**(10), 1723–1730 (2000)
32. Lysaker, M., Lundervold, A., Tai, X.-C.: Noise removal using fourth-order partial differential equation with applications to medical magnetic resonance images in space and time. *IEEE Trans. Image Process.* **12**(12), 1579–1590 (2003)
33. Didas, S., Weickert, J., Burgeth, B.: Properties of higher order nonlinear diffusion filtering. *J. Math. Imaging Vis.* **35**(3), 208–226 (2009)
34. Hajiaboli, M.R.: An anisotropic fourth-order diffusion filter for image noise removal. *Int. J. Comput. Vis.* **92**(2), 177–191 (2011)
35. Gorgi Zadeh, S., Didas, S., Wintergerst, M.W.M., Schultz, T.: Multi-scale anisotropic fourth-order diffusion improves ridge and valley localization. *J. Math. Imaging Vis.* **59**(2), 257–269 (2017)
36. Li, P., Li, S.-J., Yao, Z.-A., Zhang, Z.-J.: Two anisotropic fourth-order partial differential equations for image inpainting. *IET Image Process.* **7**(3), 260–269 (2013)
37. Jumakulyyev, I., Schultz, T.: Fourth-order anisotropic diffusion for inpainting and image compression. In: Özarslan, E., Schultz, T., Zhang, E., Fuster, A. (eds.) *Anisotropy Across Fields and Scales*, pp. 99–124. Springer, Cham (2021)
38. Chen, Y., Ranftl, R., Pock, T.: A bi-level view of inpainting-based image compression. In: Kúkelová, Z., Heller, J. (eds.) *Proc. Computer Vision Winter Workshop*, pp. 19–26 (2014)
39. Peter, P., Hoffmann, S., Nedwed, F., Hoeltgen, L., Weickert, J.: Evaluating the true potential of diffusion-based inpainting in a compression context. *Signal Process. Image Commun.* **46**, 40–53 (2016)
40. Amrani, N., Serra-Sagristà, J., Peter, P., Weickert, J.: Diffusion-based inpainting for coding remote-sensing data. *IEEE Geosci. Remote Sens. Lett.* **14**(8), 1203–1207 (2017)
41. Peter, P., Kaufhold, L., Weickert, J.: Turning diffusion-based image colorization into efficient color compression. *IEEE Trans. Image Process.* **26**(2), 860–869 (2016)
42. Beaulieu, C.: The basis of anisotropic water diffusion in the nervous system - a technical review. *NMR Biomed.* **15**(7–8), 435–455 (2002)
43. Stejskal, E.O., Tanner, J.E.: Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient. *J. Chem. Phys.* **42**(1), 288–292 (1965)
44. Pierpaoli, C., Basser, P.J.: Toward a quantitative assessment of diffusion anisotropy. *Magn. Reson. Med.* **36**, 893–906 (1996)
45. Callaghan, P.T., Eccles, C.D., Xia, Y.: NMR microscopy of dynamic displacements: k-space and q-space imaging. *J. Phys. E* **21**(8), 820–822 (1988)
46. Cheng, J., Shen, D., Yap, P.-T., Basser, P.J.: Single-and multiple-shell uniform sampling schemes for diffusion MRI using spherical codes. *IEEE Trans. Med. Imaging* **37**(1), 185–199 (2017)
47. Andersson, J.L.R., Sotiropoulos, S.N.: An integrated approach to correction for off-resonance effects and subject movement in diffusion MR imaging. *NeuroImage* **125**, 1063–1078 (2016)
48. Kassim, A.A., Yan, P., Lee, W.S., Sengupta, K.: Motion compensated lossy-to-lossless compression of 4-D medical images using integer wavelet transforms. *IEEE Trans. Inf. Technol. Biomed.* **9**(1), 132–138 (2005)
49. Sanchez, V., Nasiopoulos, P., Abugharbieh, R.: Efficient lossless compression of 4-D medical images based on the advanced video coding scheme. *IEEE Trans. Inf. Technol. Biomed.* **12**(4), 442–446 (2008)
50. Zeng, L., Jansen, C.P., Marsch, S., Unser, M., Hunziker, P.R.: Four-dimensional wavelet compression of arbitrarily sized echocardiographic data. *IEEE Trans. Med. Imaging* **21**(9), 1179–1187 (2002)
51. Lalgudi, H.G., Bilgin, A., Marcellin, M.W., Nadar, M.S.: Compression of fMRI and ultrasound images using 4D SPIHT. *IEEE Int. Conf. Image Process. (ICIP)* **2**, 746–749 (2005)
52. Liu, Y., Pearlman, W.A.: Four-dimensional wavelet compression of 4-D medical images using scalable 4-D SBHP. In: *Proc. Data Compression Conference (DCC)*, pp. 233–242 (2007). IEEE
53. Belhadef, L., Maaza, Z.M.: Lossless 4D medical images compression with motion compensation and lifting wavelet transform. *Int. J. Signal Process. Syst.* **4**(2), 168–171 (2016)
54. Nguyen, B.P., Chui, C.-K., Ong, S.-H., Chang, S.: An efficient compression scheme for 4-D medical images using hierarchical vector quantization and motion compensation. *Comput. Biol. Med.* **41**(9), 843–856 (2011)
55. Witten, I.H., Neal, R.M., Cleary, J.G.: Arithmetic coding for data compression. *Commun. ACM* **30**(6), 520–540 (1987)
56. Duda, J., Tahboub, K., Gadgil, N.J., Delp, E.J.: The use of asymmetric numeral systems as an accurate replacement for Huffman coding. In: *Proc. IEEE Picture Coding Symposium (PCS)*, pp. 65–69 (2015)
57. Merlet, S.: *Compressive sensing in diffusion MRI*. PhD thesis, Université Nice Sophia Antipolis (2013)
58. Tobisch, A., Stirnberg, R., Harms, R.L., Schultz, T., Roebroek, A., Breteler, M.M., Stöcker, T.: Compressed sensing diffusion spectrum imaging for accelerated diffusion microstructure MRI in long-term population imaging. *Front. Neurosci.* **12**, 650 (2018)
59. Nagoor, O.H., Whittle, J., Deng, J., Mora, B., Jones, M.W.: Lossless compression for volumetric medical images using deep neural network with local sampling. In: *Proc. IEEE Int'l Conf. on Image Processing (ICIP)*, pp. 2815–2819 (2020)
60. Nasr, M., Shokri, R., Houmansadr, A.: Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning. In: *Proc. IEEE Symp. on Security and Privacy*, pp. 739–753 (2019)
61. Marwood, D., Massimino, P., Covell, M., Baluja, S.: Representing images in 200 bytes: Compression via triangulation. In: *2018 25th IEEE International Conference on Image Processing (ICIP)*, pp. 405–409 (2018). IEEE
62. Peter, P.: Fast inpainting-based compression: Combining shepard interpolation with joint inpainting and prediction. In: *2019 IEEE International Conference on Image Processing (ICIP)*, pp. 3557–3561 (2019). IEEE
63. Chizhov, V., Weickert, J.: Efficient data optimisation for harmonic inpainting with finite elements. In: *Int'l Conf. on Computer Analysis of Images and Patterns*, pp. 432–441 (2021)
64. Logg, A., Mardal, K.-A., Wells, G.: *Automated Solution of Differential Equations by the Finite Element Method: The FEniCS Book*. Lecture Notes in Computational Science and Engineering, vol. 84. Springer, Berlin, Heidelberg (2012)
65. Mainberger, M., Hoffmann, S., Weickert, J., Tang, C.H., Johannsen, D., Neumann, F., Doerr, B.: Optimising spatial and tonal data for homogeneous diffusion inpainting. In: *International Conference*

- on Scale Space and Variational Methods in Computer Vision, pp. 26–37 (2011). Springer
66. Langtangen, H.P., Logg, A.: Solving PDEs in Python: The FEniCS Tutorial I, 1st edn. Springer, Oslo (2017)
  67. Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S.J., Brett, M., Wilson, J., Millman, K.J., Mayorov, N., Nelson, A.R.J., Jones, E., Kern, R., Larson, E., Carey, C.J., Polat, İ., Feng, Y., Moore, E.W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E.A., Harris, C.R., Archibald, A.M., Ribeiro, A.H., Pedregosa, F., van Mulbregt, P., SciPy 1.0 Contributors: SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods* **17**, 261–272 (2020)
  68. Jenkinson, M., Bannister, P., Brady, M., Smith, S.: Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage* **17**(2), 825–841 (2002)
  69. Leemans, A., Jones, D.K.: The B-matrix must be rotated when correcting for subject motion in DTI data. *Magn. Reson. Med.* **61**(6), 1336–1349 (2009)
  70. Garyfallidis, E., Brett, M., Amirbekian, B., Rokem, A., Van Der Walt, S., Descoteaux, M., Nimmo-Smith, I.: Dipy Contributors: Dipy, a library for the analysis of diffusion MRI data. *Front. Neuroinf.* **8**, 8 (2014)
  71. Koch, A., Zhukov, A., Stöcker, T., Groeschel, S., Schultz, T.: SHORE-based detection and imputation of dropout in diffusion MRI. *Magn. Reson. Med.* **82**(6), 2286–2298 (2019)
  72. Alexander, D.C., Barker, G.J., Arridge, S.R.: Detection and modeling of non-gaussian apparent diffusion coefficient profiles in human brain data. *Magn. Reson. Med.* **48**, 331–340 (2002)

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Thomas Schultz** is a professor at the b-it and Department of Computer Science at the University of Bonn, Germany, as well as a member of the Lamarr Institute for Machine Learning and Artificial Intelligence. He received his doctorate in 2009 from Saarland University and MPI Informatics, spent two years as a DAAD post-doctoral fellow at the University of Chicago and another two years as a post-doctoral researcher at the MPI for Intelligent Systems. His research focuses on medical image analysis and interactive visualization, with a particular focus on neuroimaging and ophthalmic imaging.



**Ikram Jumakulyyev** graduated from the Diploma program in Applied Mathematics at Turkmen State University named after Magtymguly, Turkmenistan in 2016. In 2018, he additionally received a joint M.Sc. degree in “Mathematical modeling in engineering: theory, numerics, applications” from the University of L’Aquila (Italy) and the University of Hamburg. He then joined the Visualization and Medical Image Analysis Group at the University of Bonn as a DAAD PhD fellow, working on PDE-based

lossless medical image compression.