

Super-Resolution in Diffusion-Weighted Imaging

Benoit Scherrer, Ali Gholipour, and Simon K. Warfield

Computational Radiology Laboratory, Department of Radiology
Children’s Hospital Boston, 300 Longwood Avenue, Boston, MA, 02115, USA

Abstract. Diffusion-weighted imaging (DWI) enables non-invasive investigation and characterization of the white-matter but suffers from a relatively poor resolution. In this work we propose a super-resolution reconstruction (SRR) technique based on the acquisition of multiple anisotropic orthogonal DWI scans. We address the problem of patient motions by aligning the volumes both in space and in q-space. The SRR is formulated as a maximum *a posteriori* (MAP) problem. It relies on a volume acquisition model which describes the generation of the acquired scans from the unknown high-resolution image. It enables the introduction of image priors that exploit spatial homogeneity and enables regularized solutions. We detail our resulting SRR optimization procedure and report various experiments including numerical simulations, synthetic SRR scenario and real world SRR scenario. Super-resolution reconstruction in DWI may enable DWI to be performed with unprecedented resolution.

Keywords: Diffusion imaging, super-resolution, orthogonal acquisitions.

1 Introduction

Diffusion-weighted imaging (DWI) is a key imaging technique to investigate and characterize the white-matter architecture and microstructure. It is based on the acquisition of several diffusion-sensitized images, probing the water diffusion in various directions and at various diffusion scales. DWI is, however, strongly limited by the relatively low resolution achievable by today’s imaging techniques: while individual axon diameter is on the order of $1\text{-}30\mu\text{m}$, typically achievable DWI resolution is on the order of $2\text{x}2\text{x}2\text{mm}^3$. Anisotropic acquisitions with a better *in-plane* resolution (up to $1\text{x}1\text{mm}^2$) can be performed on modern scanners but lead to a lower signal to noise ratio (SNR) and are not adapted to further perform tractography. Consequently, due to strong partial volume effect, DWI has been limited to the investigation of the major fiber ”highways” in the brain. Increasing the resolution of DWI acquisitions holds out the potential to allow investigation of novel fiber structures and will enable a more accurate white-matter and brain connectivity assessment. However, increasing the resolution is challenging in DWI. First, the common anatomical imaging resolution enhancement techniques, based on the modification of the pulse sequence, cannot be employed in DWI. The data of a same slice cannot be acquired over many excitations due to phase inconsistencies resulting from even minimal physiological

motion during the application of the sensitizing gradients. Second, the SNR is directly proportional to the voxel size, and proportional to the square root of the number of averages. Consequently, 64 averages are necessary to increase the resolution from $2\text{x}2\text{x}2\text{mm}^3$ to $1\text{x}1\text{x}1\text{mm}^3$ while ensuring a similar SNR. A 5 minute acquisition would become a 5 hour scan, which is not realistic.

Solutions to achieve higher resolution include improvements of the MRI scanner hardware itself, such as employing higher magnetic fields (7 Tesla, 11 Tesla) or stronger and faster gradients. Another solution is to consider adapted acquisitions and post-processing algorithms. Super-resolution (SR) approaches were originally developed for the reconstruction of high-resolution (HR) images from a set of low-resolution (LR) images in video sequences [6]. SR techniques have also been applied to anatomical magnetic resonance imaging (MRI) [4]. Calamante *et al.* [1] have explored the application of interpolation of fiber tracts inside voxels but do not increase the resolution of the imaging data. To our knowledge, only [9] have used SR in DWI. They proposed to employ the Irani-Peleg SR technique [6] from a set of spatially in-plane subpixel-shifted scans. However, MRI being a Fourier acquisition technique, employing in-plane shifting has been shown to be equivalent to a global phase shift in k-space [3]. Such a technique does not enable any resolution enhancement in MRI but is equivalent to interpolation by zero-padding of the raw data in the temporal domain. Recently, sub-voxel spatial shifts in the *slice-select* dimension have been shown to enable SR in anatomical MRI [4]. Scattered data interpolation has been used to combine multiple DWI images of moving subjects [7]. Other techniques using multiple and *orthogonal* fast slice scans have enabled the SR reconstruction of moving subjects in fetal imaging [2].

In this work we propose to investigate a novel super-resolution reconstruction (SRR) approach for DWI. It is based on the acquisition of *multiple anisotropic orthogonal* DWI scans (see Fig.1a). First, we propose a technique to align each volume both in space and in q-space. Second, we formulate the super-resolution reconstruction from multiple scans as a maximum *a posteriori* estimation problem. Inspired by recent developments in fetus anatomical imaging [2], our approach relies on an image acquisition model. It describes the generation of the acquired volumes from the unknown HR volumes we aim to recover. Our formulation enables introduction of image priors that exploit spatial homogeneity and provide regularized solutions. We report various experiments including numerical simulations, synthetic SRR scenario and real world SRR scenario. The results indicate resolution enhancement in DWI through SRR.

2 Material and Methods

DW signal smoothness hypothesis and interpolation in q-space. We consider K *orthogonal* DWI LR acquisitions containing G sensitizing-gradients each. The KG volumes should all be properly aligned to enable the super-resolution reconstruction. First we perform an alignment in space: we register each volume to a reference volume, chosen as the $B = 0\text{s/mm}^2$ volume of the first DWI LR




Fig. 1. (a) Scheme illustrating the super-resolution reconstruction from the acquisition of two orthogonal thick slices. (b) Alignment in q-space: the gradient images of each acquisition $k > 1$ are resampled so that its gradient directions $\mathbf{g}^k$ (red dots arrows) correspond to the reference gradient directions of the first acquisition $\tilde{\mathbf{g}}$ (grey plain arrows). At each voxel, we compute novel intensities corresponding to the gradients $\tilde{\mathbf{g}}$ by interpolation in q-space from the observed intensities corresponding to $\mathbf{g}^k$.

acquisition. Each gradient orientation is compensated for the rotation component of the transformation, providing a gradient set $\mathbf{g}^k = (g_1^k, \dots, g_G^k)$ for each acquisition k .

Because of possible patient motions between the scans, the resulting G gradient images may correspond to slightly modified sensitizing gradients across the K DWI acquisitions. However, it is *essential* that the images combined by the SRR technique correspond to the same “scene”, namely that they correspond to the exact same gradient direction and show identical diffusion-attenuation patterns. Consequently, we propose to perform an alignment of the volumes in q-space (see Fig.1b). We consider that the DW-signal varies smoothly in q-space and propose to resample the gradient images. We consider the gradients of the first DWI LR acquisition as the reference gradients $\tilde{\mathbf{g}} = \mathbf{g}^1$. We align in q-space each other $k \geq 2$ DWI acquisition so that their gradients exactly match $\tilde{\mathbf{g}}$. This is done by using interpolation in q-space. At each voxel, we consider the G intensity values corresponding to the gradients $\mathbf{g}^k$. We then compute the new intensity values corresponding to the gradients $\tilde{\mathbf{g}}$. The interpolation is performed via Kriging [8], a general and efficient statistical interpolation framework originally introduced for geology and mining applications. This enables us to easily perform scattered data interpolation. It determines the weights of the contribution of each observed data via the resolution of a simple linear system. In the absence of motion (*i.e.* a gradient g_j exactly matches a gradient $\tilde{g}_{j'}$), the interpolated intensity match exactly the observed intensity. As a result, the K LR acquisitions are all aligned in space and represent the same sensitizing-gradient set $\tilde{\mathbf{g}}$.

Super-resolution model-based reconstruction. In this section each LR volume is considered to be aligned in space and in q-space. The SRR technique we now address is performed for each gradient image separately. For each gradient $\tilde{g}_j$, we aim to recover the HR image $\mathbf{x}_j$ underlying the K LR images $\mathbf{y}_j = (\mathbf{y}_{1,j}, \dots, \mathbf{y}_{K,j})$. By omitting the gradient dependency to simplify the notations, we consider the K LR volumes $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_K)$ to be the degraded version of the same unknown HR volume $\mathbf{x}$. $\mathbf{x}$ is estimated according to the

maximum *a posteriori* principle, by maximizing:

$$\hat{\mathbf{x}}_{\text{MAP}} = \arg \max_{\mathbf{x}} p(\mathbf{x}|\mathbf{y}) = \arg \max_{\mathbf{x}} p(\mathbf{y}|\mathbf{x})p(\mathbf{x}) = \arg \max_{\mathbf{x}} [\ln p(\mathbf{y}|\mathbf{x}) + \ln p(\mathbf{x})] , \quad (1)$$

which decomposes into a likelihood term and a prior term.

Likelihood term $p(\mathbf{y}|\mathbf{x})$: we consider a model describing how the LR volumes are obtained from the unknown HR volume. In a trade-off between a realistic model and a feasible solution, we consider the following acquisition model:

$$\mathbf{y}_k = \mathbf{D}_k \mathbf{B}_k \mathbf{M}_k \mathbf{x} + \epsilon_k , \quad (2)$$

where the volumes $\mathbf{y}_k$ and $\mathbf{x}$ are expressed as column vectors by a lexicographical reordering of the pixels, $\mathbf{D}_k$ is a down-sampling matrix, $\mathbf{M}_k$ is the warping matrix that maps the HR volume $\mathbf{x}$ to the LR volume $\mathbf{y}_k$ and ϵ_k is the residual noise vector. $\mathbf{B}_k$ is the blur, or point spread function (PSF) of the MRI signal acquisition process. It is constructed from the imaging parameters. The PSF can be separated into three components corresponding to the slice-selection direction and the phase- and frequency-encoding directions. As in [3], we currently consider a PSF in the slice-selection direction only, which describes the slice selection profile. We consider a Gaussian slice selection profile of variance σ_{PSF}^2 . Consequently, on the basis of Eq.2, the unknown HR volume $\mathbf{x}$ goes through geometric and signal operations, including motion, signal averaging, and resampling, to generate the acquired LR volume $\mathbf{y}_k$. Assuming a Gaussian noise with zero-mean and variance σ_k for ϵ_k , the likelihood of the LR volume $\mathbf{y}_k$ under the model in Eq.2 can be written as:

$$p(\mathbf{y}_k|\mathbf{x}, \sigma_k) = \frac{1}{\sigma_k \sqrt{2\pi}} \exp \left(-\frac{\|\mathbf{y}_k - \mathbf{D}_k \mathbf{B}_k \mathbf{M}_k \mathbf{x}\|^2}{2\sigma_k^2} \right) . \quad (3)$$

Assuming statistical independence of the noise between the acquisitions, we have $p(\mathbf{y}|\mathbf{x}, \sigma) = \prod_{k=1}^K p(\mathbf{y}_k|\mathbf{x}, \sigma_k)$ with $\sigma = (\sigma_1, \dots, \sigma_K)$.

Prior term $p(\mathbf{x})$: the term $p(\mathbf{x})$ in Eq.1 enables us to incorporate a prior knowledge on $\mathbf{x}$. In this work we consider a regularization prior that exploits spatial homogeneity. More precisely, we favour smoothness of $\mathbf{x}$ by setting $p(\mathbf{x}|\lambda) = \exp(-\lambda \|\mathbf{Q}\mathbf{x}\|^2)$ where the matrix $\mathbf{Q}$ is symmetric definite positive and represents a linear high-pass operation. The parameter λ controls the regularization strength. In this work, $\mathbf{Q}$ is chosen as the 3-D discrete Laplacian corresponding to the following approximation of the partial derivative for a 3-D image $\mathbf{I}$ indexed by $\mathbf{u} \in \mathbb{N}^3$, and for a direction $\mathbf{u}_m \in \mathbb{N}^3$ ($m \in \{1, 2, 3\}$):

$$\partial_m \mathbf{I}(\mathbf{u}) \approx \left(\mathbf{I}(\mathbf{u} + \mathbf{u}_m) - 2\mathbf{I}(\mathbf{u}) + \mathbf{I}(\mathbf{u} - \mathbf{u}_m) \right) / (2\|\mathbf{u}_m\|) . \quad (4)$$

Ultimately, by considering the same σ_k across acquisitions, maximization of the posterior distribution in Eq.1 leads to the following minimization:

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \sum_{k=1}^K \|\mathbf{y}_k - \mathbf{D}_k \mathbf{B}_k \mathbf{M}_k \mathbf{x}\|^2 + \lambda \|\mathbf{Q}\mathbf{x}\|^2 . \quad (5)$$

DWI-SRR optimization procedure. The matrix $\mathbf{D}_k \mathbf{B}_k \mathbf{M}_k$ is especially large and the classical solution through pseudo-inverse is prohibitive. Instead we use a steepest descent iterative minimization approach. Differentiation of Eq.5 leads to the following update at each step:

$$\hat{\mathbf{x}}^{n+1} = \hat{\mathbf{x}}^n - \alpha \left[\sum_{k=1}^K \mathbf{M}_k^T \mathbf{B}_k^T \mathbf{D}_k^T (\mathbf{D}_k \mathbf{B}_k \mathbf{M}_k \hat{\mathbf{x}}^n - \mathbf{y}_k) - \lambda \mathbf{Q}^T \mathbf{Q} \hat{\mathbf{x}}^n \right], \quad (6)$$

where α is the step size and $\mathbf{M}_k^T$ denotes the transpose of $\mathbf{M}_k$. The iterative algorithm is initialized by setting $\hat{\mathbf{x}}^0$ equal to the mean of the aligned LR volumes. The iterative minimization is stopped when $\|\hat{\mathbf{x}}^{n+1} - \hat{\mathbf{x}}^n\|_1 < \mathcal{T}_{\text{STOP}}$. The DWI-SRR optimization procedure is synthesized by the following *pseudo-code*:

```

FOR each LR study  $k$ 
  FOR each gradient image of  $k$ 
     $\mathbf{M}_k^T \leftarrow$  Register to the reference volume
    Apply the transform to the gradient
  ENDFOR
ENDFOR
FOR each output gradient  $\mathbf{g}_g$ 
  FOR each LR study  $k$ 
     $\mathbf{y}_k \leftarrow$  Compute the gradient image for  $\mathbf{g}_g$  (q-space interpolation)
  ENDFOR
  Compute  $\hat{\mathbf{x}}^0 \leftarrow$  Mean of the  $\mathbf{M}_k^T \mathbf{D}_k^T \mathbf{y}_k$ 
  WHILE  $\|\hat{\mathbf{x}}^n - \hat{\mathbf{x}}^{n-1}\|_1 \geq \mathcal{T}_{\text{STOP}}$ 
     $\hat{\mathbf{x}}^{n+1} \leftarrow$  Update with Eq. 6
  ENDWHILE
ENDFOR

```

3 Results

The SRR procedure was implemented in C++ and optimized with various techniques to reduce the processing burden. The $\mathbf{M}_k^T \mathbf{B}_k^T \mathbf{D}_k^T \mathbf{D}_k \mathbf{B}_k \mathbf{M}_k$ and $\mathbf{M}_k^T \mathbf{B}_k^T \mathbf{D}_k^T \mathbf{y}_k$ matrices were precomputed, and the derivative of the Laplacian corresponding to the finite difference scheme in Eq.4 computed analytically. To accelerate the convergence, the steepest descent algorithm was implemented with a variable per-voxel step-size α which incorporates inertia: the step-size is multiplied by 1.1 when the sign of two consecutive computed gradient does not change, and divided by two otherwise. α is initialized to 0.01 and constrained to lie in $[0.1, 10^{-6}]$. The FWHM for the Gaussian slice model was set to half the slice thickness, by setting $\sigma_{\text{PSF}} = (\text{slice thickness}) / (4\sqrt{2} \ln 2)$. Other parameters were set to $\lambda = 0.001$ and $\mathcal{T}_{\text{STOP}} = 10^{-5}$.

Numerical simulations. We first performed numerical simulations. The DW-signal for tensors (see Fig.2a) was simulated with a b-value of 1000s/mm² for 15 directions, and corrupted by Rician noise (SNR of 30dB for the b=0s/mm² image). Linear down-sampling with a factor of 4 was applied to each gradient image in each of the three directions, providing three simulated LR acquisitions. Various tensor estimation were performed (see Fig.2b-d), and the corresponding fractional anisotropy (FA) computed (Fig.2e-g) (ground truth FA=0.8).

Synthetic SRR scenario. We then simulated a SRR scenario by down-sampling




Fig. 2. Numerical simulations from the tensors of Fig.a. Fig.b-d: Tensors estimated resp. from a single LR acquisition, from the mean of the LR acquisitions and from the SRR. Fig.e-g: Corresponding tensor fractional anisotropy. It shows the tensor directions to be well estimated from the mean (Fig.c). However, the SRR provides a much more accurate reconstruction of the complete tensor (see the better FA uniformity in Fig.g).




Fig. 3. Fig.a: Synthetic SRR scenario from a real acquisition. a.a: $b=0$ image. a.b: Axial down-sampled $b=0$ image with a factor of 4. a.c: Mean of the $b=0$ images of the LR acquisitions. a.d: SRR of the $b=0$ image. The SRR is better contrasted and is less blurry than the mean. Fig.b and Fig.c: Quantitative evaluation of the reconstruction accuracy in term of PSNR for the $1/2$ and $1/4$ down-sampling, for *each* gradient direction.

in each of the three directions a real DWI acquisition (Siemens 3T Trio, 32 channel head coil, 68 slices, FOV=220mm, matrix=128x128, resolution=1.7x1.7x2mm³, TE=86ms/TR=8800ms, 30 directions at $B=1000\text{s/mm}^2$, 5 $B=0\text{s/mm}^2$). A down-sampling of factor 2 and 4 were considered. The SRR at the original resolution was estimated and qualitatively compared to the original image (see Fig.3a). The SRR estimation time was approximately 2 hours on a 3Ghz Intel Xeon (3 to 4min per gradient image). Fig.3b-c report, for each gradient direction, the Peak Signal to Noise Ratio (PSNR) with the original acquisition. The PSNR is defined by $20 \log_{10}(\text{MAX}/\sqrt{\text{MSE}})$ with MAX the maximum intensity and MSE the mean square error. Fig.4 synthesizes, for four voxels, the error with the ground-truth for all the gradient directions via a 3-D angular reconstruction.

Real SRR scenario. Finally, we performed the acquisition of three anisotropic DWI scans (same parameters as before, except: 1.6x1.6mm² in-plane res., 5mm slice




Fig. 4. 3-Dimensional angular reconstructions of the diffusion signal at four voxels whose position is shown on the $b=0\text{s/mm}^2$ image (left image). The voxels were chosen to have a high FA ($FA > 0.9$). The obtained 3-D shapes are proportional to the apparent diffusion coefficient (ADC). Comparison between the 3-D reconstruction performed from the mean image (first line) and from the SRR estimate (second line). The stick indicates the major fiber direction estimated by a single-tensor model. The color encodes for the error with the ground-truth (difference in image intensity). It shows the SRR estimate to provide a much better reconstruction for each gradient image.




Fig. 5. Real SRR scenario. Fig.a-b: FA computed from the mean of the LR acquisition (a) and from the SRR (b). Fig.c-d: idem for MD. It shows that the SRR leads to more contrasted and less blurry FA and MD estimates.

thickness, 38 slices, $TE/TR=87/4700\text{ms}$) and achieved the reconstruction at $1.6 \times 1.6 \times 2.5\text{mm}^3$. Fig.5 shows the FA and the mean-diffusivity (MD) computed from the mean of the LR acquisitions and from our SRR technique. As in Fig.2, the images are less blurry and more contrasted when employing our SRR technique.

4 Discussion

We have proposed a novel SRR technique for DWI based on the acquisition of orthogonal anisotropic DWI scans. To our knowledge, it is the first attempt to perform SRR in DWI in the last decade, since [9]. In contrast to [9], we take into account possible patient motions by aligning the volumes in both space *and* q-space. In addition, we formulate the SRR as a MAP estimation problem. For each gradient, we estimate the underlying unknown HR volume given the acquired LR DWI scans. Our formulation enables us to integrate an image acquisition model and to integrate image priors. We have shown that the SRR estimate outperforms the mean of the LR acquisitions, providing a better contrast and less blurry

results for the gradient images (Fig.3a) and for diffusion parameters such as the FA or the MD (Fig.2 and Fig.5). The quantitative evaluation showed an increase of PSNR on the order of 6dB and 2dB for our two synthetic down-sampling scenarios (Fig.3b-c). In future work, we will evaluate the benefits of correcting for the geometric distortions. Indeed, the acquisitions show locally very different geometric distortion patterns due to different phase-encoding directions. As a result, accurate alignment of the images is difficult, which can locally perturb the SRR. We will investigate the effectiveness of employing a distortion correction technique by acquisition of a magnetic field map [10] or by the acquisition of two DWI scans with reversed phase directions [5]. Finally, we will investigate the introduction of novel priors. Particularly, incorporating the brain anatomy description provided by a HR T2-w scan may enable improved SRR in DWI.

SRR techniques are of great interest for medical imaging because they enable us to go beyond the limits dictated by the hardware. With today's scanners, they may enable routine HR investigation of the brain white-matter in clinically compatible scan time. Combined with future MRI hardware improvements, they may enable DWI to be performed with unprecedented resolution.

References

1. Calamante, F., Tournier, J.D., Jackson, G.D., Connelly, A.: Track-density imaging (TDI): super-resolution white matter imaging using whole-brain track-density mapping. *NeuroImage* 53(4), 1233–43 (2010)
2. Gholipour, A., Estroff, J.A., Warfield, S.K.: Robust super-resolution volume reconstruction from slice acquisitions: application to fetal brain MRI. *IEEE Trans. on Med. Imag.* 29(10), 1739–58 (2010)
3. Greenspan, H.: MRI inter-slice reconstruction using super-resolution. *Magn. Res. Imag.* 20(5), 437–446 (2002)
4. Greenspan, H.: Super-Resolution in Medical Imaging. *Comput. J.* 52(1), 43–63 (2009)
5. Holland, D., Kuperman, J.M., Dale, A.M.: Efficient correction of inhomogeneous static magnetic field-induced distortion in echo planar imaging. *NeuroImage* 50(1), 175–83 (2010)
6. Irani, M., Peleg, S.: Motion Analysis for Image Enhancement : Resolution, Occlusion, and Transparency. *J. Vis. Com. and Image Rep.* 4(4), 324–335 (1993)
7. Jiang, S., Xue, H., Counsell, S., Anjari, M., Allsop, J., Rutherford, M., Rueckert, D., Hajnal, J.V.: Diffusion tensor imaging (DTI) of the brain in moving subjects: application to in-utero fetal and ex-utero studies. *Magn. Res. in Med.* 62(3), 645–55 (2009)
8. Matheron, G.: Principles of geostatistics. *Economic geology* 58(8), 1246–66 (1963)
9. Peled, S., Yeshurun, Y.: Superresolution in MRI: application to human white matter fiber tract visualization by diffusion tensor imaging. *Magn. Res. in Med.* 45(1), 29–35 (2001)
10. Windischberger, C., Robinson, S., Rauscher, A., Barth, M., Moser, E.: Robust field map generation using a triple-echo acquisition. *J. Magn. Reson. Imaging* 20(4), 730–734 (2004)

This investigation was supported in part by NIH grants R01 RR021885, R01 EB008015, R03 EB008680, R01 LM010033 and UL1 RR025758-03.