Accuracy and precision in super-resolution MRI: Enabling spherical tensor diffusion encoding at ultra-high b-values and high resolution

Geraline Vis[a], **Markus Nilsson**[a], **Carl-Fredrik Westin**[b], **Filip Szczepankiewicz**[a,b,*]

[a]Department of Diagnostic Radiology, Clinical Sciences Lund, Lund University, Lund, Sweden

[b]Department of Radiology, Brigham and Women's Hospital Harvard Medical School, Boston, MA, United States

Abstract

Diffusion MRI (dMRI) can probe the tissue microstructure but suffers from low signal-to-noise ratio (SNR) whenever high resolution is combined with high diffusion encoding strengths. Low SNR leads to poor precision as well as poor accuracy of the diffusion-weighted signal; the latter is caused by the rectified noise floor and can be observed as a positive bias in magnitude signal. Super-resolution techniques may facilitate a beneficial tradeoff between bias and resolution by allowing acquisition at low spatial resolution and high SNR, whereafter high spatial resolution is recovered by image reconstruction. In this work, we describe a super-resolution reconstruction framework for dMRI and investigate its performance with respect to signal accuracy and precision. Using phantom experiments and numerical simulations, we show that the super-resolution approach improves accuracy by facilitating a more beneficial trade-off between spatial resolution and diffusion encoding strength before the noise floor affects the signal. By contrast, precision is shown to have a less straightforward dependency on acquisition, reconstruction, and intrinsic tissue parameters. Indeed, we find a gain in precision from super-resolution reconstruction is substantial only when some spatial resolution is sacrificed. Finally, we deployed super-resolution reconstruction in a healthy brain for the challenging combination of spherical b-tensor encoding at ultra-high b-values and high spatial resolution—a configuration that produces a unique contrast that emphasizes tissue in which diffusion is restricted in all directions. This demonstration showcased that super-resolution reconstruction enables a vastly superior image contrast compared to conventional imaging, facilitating investigations that would otherwise have prohibitively low SNR, resolution or require non-conventional MRI hardware.

Keywords

1. Introduction

Diffusion MRI (dMRI) is a non-invasive method for investigating tissue microstructure in healthy and pathological tissue. Investigations of subtle microstructure features rely on the use of strong diffusion-weighting (ultra-high b-values) or tensor-valued diffusion encoding (Jensen et al., 2005; Tax et al., 2020; Westin et al., 2016), both of which frequently suffer from low signal-to-noise ratios (SNR). Low SNR may cause poor signal accuracy in magnitude imaging due to the so-called rectified noise floor, which causes a positive signal bias (Gudbjartsson and Patz, 1995). Unlike precision, signal accuracy is not improved by averaging over magnitude signals unless the phase is coherent across the observations (Jones and Basser, 2004). This is challenging because the diffusion encoding causes phase variation in the presence of tissue motion (Eichner et al., 2015; Prah et al., 2010). Therefore, it is challenging to perform accurate measurements of diffusion-weighted signal at low SNR.

In addition to strong diffusion-weighting, an approach that regularly suffers from low SNR is spherical or isotropic b-tensor encoding, especially at ultra-high b-values. Although high b-value spherical encoding is challenging, it is desirable because it emphasizes signal from tissue in which diffusion is restricted along all directions. For example, it can be used to highlight the tightly packed granule cells in the cerebellar cortex, which are affected in diseases such as spinocerebellar ataxis and Alzheimer disease (Fukutani et al., 1996; Xia et al., 2013). So far, this contrast has been demonstrated at preclinical MRI (Lundell et al., 2019) and MRI systems with ultra-strong gradients and low spatial resolution (Tax et al., 2020). Making this imaging contrast available at clinical MRI scanners and high spatial resolution may enable studies of the cerebellum in a wide range of neurological conditions.

Super-resolution reconstruction (SRR) is a promising solution to the problem of low SNR in dMRI. In SRR, data is acquired at a low spatial resolution—with improved precision and accuracy—and a specialized image reconstruction is used to reconstruct a high-resolution image. Several approaches have been proposed throughout the years, showing SRR methods can balance the trade-off between SNR, spatial resolution and acquisition time (Greenspan, 2009; Greenspan et al., 2002; Jeurissen et al., 2018; Plenge et al., 2012; Poot et al., 2013; Scherrer et al., 2012; Shilling et al., 2009; Souza and Senn, 2008; Van Reeth et al., 2012; Van Steenkiste et al., 2016; Yang et al., 2018). For example, Yang et al. proposed a strategy where joint information from the adjacent scanning directions was used to improve resolution. Poot et al. (2012) demonstrated increased resolution of diffusion tensor parameters from a set of super-resolved diffusion-weighted images with fully sampled diffusion-weightings and gradient directions, while Van Steenkiste et al. (2016) showed increased spatial resolution of diffusion tensor parameters when optimizing both k- and q-space sampling. More recently, Jeurissen et al. (2018) showed an improvement of both accuracy and precision in q-space trajectory imaging parameter estimation, where higher diffusion-weightings are used. These studies also illustrate a flexibility in how the low-resolution data is acquired. For example, the image acquisition can comprise three orthogonal low-resolution axes (Souza and Senn, 2008), slice shifting along the low-resolution direction (Greenspan et al., 2002), or multiple stacks of slices rotated about a common axis (Shilling et al., 2009). This flexibility allows SRR to be tailored to the needs

of echo-planar imaging (EPI) which is the imaging readout strategy most used in dMRI. For example, using a common phase encoding direction across the low-resolution acquisitions avoid that local field inhomogeneities cause variable geometric distortions (Jezzard et al., 1998).

Although SRR has seen uptake in several MRI applications, the literature lacks a systematic treatment of noise propagation and how it affects accuracy and precision. For example, imaging at a low resolution allows faster imaging (shorter repetition times), but leads to a complex interplay between SNR, scan time, repetition time, and T1-relaxation. An understanding of these trade-offs is necessary for an optimized experimental design. Furthermore, the application of SRR for spherical b-tensor encoding at ultra-high b-values and high resolution are yet to be explored. In this work, we aim to formally analyze noise propagation in SRR and experimentally verify the impact of SRR on signal accuracy and precision. Finally, we demonstrate the value of SRR in the challenging combination of spherical b-tensor encoding at a b-value of 4.0 ms/μm^2 and 1.6 mm^3 isotropic resolution in the brain.

2. Theory

Super-resolution reconstruction aims at recovering a high-resolution image from multiple low-resolution images that sample the object in complementary ways. Here we acquired multiple stacks of thick slices rotated around the phase encoding direction (Shilling et al., 2009). To reconstruct an image with isotropic voxels, the lower limit of low-resolution rotations N_R is given by (Plenge et al., 2012)

$$N_R \geq \frac{\pi}{2} \cdot \alpha ,$$ (1)

where α is the aspect factor, i.e., the ratio between the long and short axis of the voxel. Note that this factor also captures the volume ratio between the low and high-resolution voxels, assuming that all other parameters are unchanged, and that the SNR is proportional to the voxel volume in multi-slice acquisitions (Edelstein et al., 1986).

The mapping from a high-resolution image ($\mathbf{x}$) to a low resolution image ($\mathbf{y}$) can be described by a linear system (Elad and Feuer, 1997)

$$\mathbf{y}_k = \mathbf{A}_k \mathbf{x} + \varepsilon_k,$$ (2)

where k is an index of the low-resolution image sample and ε is random noise. The sampling matrix $\mathbf{A}_k$ describes the rotation/translation, down-sampling and blurring of the high-resolution image. As these factors are known, $\mathbf{A}_k$ can be constructed from the image dimensions, voxel spacing and slice profile. Here, we used a box function to describe the slice profile. Both $\mathbf{y}_k$ and $\mathbf{x}$ are expressed as column vectors such that a square image with N voxels on each side is represented by a $N^2 \times 1$ vector. The complete sampling of all low-resolution images can be described in a single linear system

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \varepsilon. \tag{3}$$

Assuming the noise is independent and normally distributed with zero mean, the solution to recover $\mathbf{x}$ given $\mathbf{A}$ and $\mathbf{y}$ can be expressed as a least-squares problem, such that

$$\widehat{\mathbf{x}} = \underset{\mathbf{x}}{\mathrm{argmin}} \parallel \mathbf{A}\mathbf{x} - \mathbf{y} \parallel_2^2. \tag{4}$$

The greater the number of complementary observations in $\mathbf{y}$, the better the condition of the problem becomes. However, the problem in Eq. (4) often remains ill-posed due to the down-sampling operation included in $\mathbf{A}$. There are several ways of solving the reconstruction problem, as thoroughly described in (Plenge et al., 2012). The solution requires some form of regularization, which often translates to imposing a smoothness to the solution (Van Reeth et al., 2012). A common approach is Tikhonov regularization (Tikhonov and Arsenin, 1977), that penalizes high spatial-frequencies in the estimated high-resolution image. Including this regularization, the regularized least-squares problem becomes

$$\widehat{\mathbf{x}} = \underset{\mathbf{x}}{\mathrm{argmin}} \parallel \mathbf{A}\mathbf{x} - \mathbf{y} \parallel_2^2 + \lambda \parallel \mathbf{R}(\mathbf{x}) \parallel_2^2, \tag{5}$$

where $\mathbf{R}$ is the regularization term and λ is a scalar weight. We will here use a general regularization term independent of the image content: $\mathbf{R}(\mathbf{x}) = \mathbf{I}$, where $\mathbf{I}$ is the identity matrix. This enables the use of the closed form solution, according to

$$\widehat{\mathbf{x}} = \left(\mathbf{A}^T\mathbf{A} + \lambda \mathbf{I}\right)^{-1}\mathbf{A}^T\mathbf{y}. \tag{6}$$

Note that $\mathbf{A}^T\mathbf{y}$ produces the average low-resolution signal on the high-resolution grid, i.e., a blurred image. Without regularization ($\lambda = 0$), the remaining term $(\mathbf{A}^T\mathbf{A})^{-1}$ is a sharpening operator, ideally reproducing the true image when applied to $\mathbf{A}^T\mathbf{y}$. As λ increases, the sharpening is reduced. However, as λ alters the denominator in Eq. (6), the intensity scale of $\hat{x}$ is dependent on λ. To remove this dependence, and thereby simplify comparisons among sampling schemes as described later, we can rewrite Eq. (6) according to

$$\widehat{\mathbf{x}}' = \left((1 - \lambda')\mathbf{A}^T\mathbf{A} + \lambda' \cdot N_R \cdot \alpha \cdot \mathbf{I}\right)^{-1}\mathbf{A}^T\mathbf{y} = \mathbf{C}\mathbf{y}, \tag{7}$$

where $\mathbf{C}$ contains the entire reconstruction operation. The new regularization factor is constrained to $0 \leq \lambda' \leq 1$, such that $\lambda' = 1$ merely returns the average low-resolution signal on the high-resolution grid, but corrected for the intensity gain caused by larger voxel volumes. Fig. 1 illustrates the SRR process for an in vivo acquisition for weak, moderate, and strong levels of regularization (different values of λ'). Weak regularization amplifies noise, while strong regularization results in a blurred image. Moderate regularization balances the two.

2.1. Signal accuracy and precision

As we evaluate the performance of SRR in terms of signal accuracy and signal precision, we will outline the tools used to assess these. Across repeated measurements under identical conditions, signal *accuracy* (or trueness)[1] concerns the distance between the average signal to the true value, while signal *precision* concerns the spread of the signal. Both terms are influenced by the distribution that characterizes the MR signal. The noise in the complex MR signal is normally distributed, whereas the magnitude signal used in dMRI is approximately Rice distributed (Gudbjartsson and Patz, 1995) (for a detailed review of MR data distributions, see (den Dekker and Sijbers, 2014) or (Veraart et al., 2013)). Consequently, an approximation of the measured signal M is given by (Gudbjartsson and Patz, 1995)

$$M = \sqrt{S^2 + \left(\sigma\sqrt{\pi/2}\right)^2}, \tag{8}$$

where S is the signal in the absence of noise and σ is the standard deviation of the noise. In absence of a true signal, the rectified noise floor causes the expected value to be $\sqrt{\pi/2}$. This gives rise to a positive signal bias that becomes prominent at low SNR. Therefore, we employ the signal-to-noise-floor ratio (SNFR) as a measure of overall signal accuracy, defined as (Lätt et al., 2007)

$$\text{SNFR} = \frac{\bar{S}(0)}{\sigma\sqrt{\pi/2}}, \tag{9}$$

where $\bar{S}(0)$ is the mean of measured non-diffusion weighted signal over repeated measurements. Note that the SNFR characterizes the maximum achievable signal attenuation for accurate signal sampling. As a measure of precision, we use the signal-to-noise ratio (SNR) defined as

$$\text{SNR} = \frac{\bar{S}(0)}{\sigma}. \tag{10}$$

Note that SNR and SNFR are both proportional to the voxel volume whereas noise levels are independent of the voxel volume (Edelstein et al., 1986).

To enable comparisons of precision between SRR and a conventional high-resolution acquisition (referred to as direct sampling), we define the SNR efficiency factor (ρ), according to

$$\rho = \frac{\text{SNR}_{\text{SRR}}}{\text{SNR}_{\text{D}}} \cdot \sqrt{\frac{\text{TA}_{\text{D}}}{\text{TA}_{\text{SRR}}}}, \tag{11}$$

---

[1]We use the terminology where accuracy is only a description of systematic errors. This is referred to as 'trueness' in the more recent ISO standard (BS ISO 5725-1, 1994).

where SNR_{SRR} and SNR_D are the SNR levels and TA_{SRR} and TA_D the acquisition times of SRR and direct sampling at a given spatial resolution. To evaluate Eq. (11), we study how noise propagates from **y** into **x**. For SNR > 3 the noise distribution is approximately Gaussian, independent and identically distributed (Cárdenas-Blanco et al., 2008; Gudbjartsson and Patz, 1995). Under the assumption of identical noise levels among the individual low-resolution images, the signal variance in the reconstructed image can be easily computed from the linear operations in Eq. (7) based on the additive property of variance. We define the noise propagation factor (κ) as the average ratio of the noise standard deviation in the high-resolution reconstructed and low-resolution images, according to

$$\kappa = \frac{1}{n} \sum_{i=1}^{n} \frac{\sigma_x(i)}{\sigma_y} = \frac{1}{n} \sum_{i=1}^{n} \sqrt{\sum_{j=1}^{m} \mathbf{C}(i,j)^2} \tag{12}$$

where n is the total number of reconstructed voxels in the reconstructed image, indexed by i, and m is the total number of low-resolution input voxels, indexed by j. The noise propagation factor is dependent on properties of the sampling (**A**, N_R, α) and the regularization (λ').

To characterize the impact of SRR on SNR efficiency, we also consider how the signal level changes between a directly sampled and an SRR acquisition. For a setup where the only difference between the two is related to the SRR configuration (slice thickness, number of slices, repetition time and number of slice orientations), Eq. (11) can be extended to include the relevant effects of imaging parameters (Appendix A) to

$$\rho = \frac{1}{\kappa} \cdot \frac{1 - \exp(-TR_{SRR}/T1)}{1 - \exp(-TR_D/T1)} \cdot \sqrt{\frac{TR_D}{N_R \cdot TR_{SRR}}}, \tag{13}$$

where TR_D and TR_{SRR} are the repetition times for direct sampling and SRR respectively, and T1 is the longitudinal relaxation time. The TR is approximately proportional to the number of slices, and the number of slices to cover the same area in an SRR acquisition can be reduced with a factor of α. Therefore, the minimal TR_{SRR} is given by

$$TR_{SRR} = \frac{TR_D}{\alpha}. \tag{14}$$

When TR_D $TR_D \gg T1$, both $1 - \exp(-TR_{SRR}/T1)$ and $1 - \exp(-TR_D/T1)$ will approach unity, as will their ratio. When sampling directly, approximately 1% signal is lost due to incomplete T1 recovery when $TR_D/T1 = 4.6$. However, when using $\alpha = 6$ and a minimal TR_{SRR}, this corresponds to a signal loss of 46%. For this reason, as the aspect factor increases and the TR is minimized, the expected SRR gain will eventually be negated by incomplete T1-recovery.

3. Methods

We evaluated the performance of the SRR by three lines of investigation. First, we investigated signal accuracy in a water phantom, both numerically and experimentally. Second, we investigated signal precision in terms of SNR efficiency, both numerically, analytically, and experimentally in vivo. Third, we demonstrated the utility of SRR for spherical b-tensor dMRI at ultra-high b-values in vivo. All simulations and data analysis were performed in Matlab (The MathWorks, Inc., Natick, Massachusetts, USA). The SRR framework is available on https://github.com/filip-szczepankiewicz/Vis_NIMG_2021/.

3.1. Data acquisition

All experiments were performed on a 3T-scanner (MAGNETOM Prisma, Siemens Healthcare, Erlangen, Germany) using a 20-channel head and neck coil. The study was approved by the local ethics committee and informed consent was obtained from all volunteers. A prototype pulse sequence was implemented based on a single-shot spin-echo sequence with echoplanar imaging readout that facilitates user-defined gradient waveforms for diffusion encoding (Szczepankiewicz et al., 2019a). Gradient waveforms for spherical b-tensor encoding were numerically optimized (Sjölund et al., 2015), including compensation for concomitant gradients (Szczepankiewicz et al., 2019b). The waveforms were constrained to a maximal gradient magnitude of 80 mT/m (L2-norm) and a slew rate of 100 T/m/s (Appendix B). Detailed information on acquisition parameters will be described per experiment.

Low-resolution data was acquired with slices rotating around a fixed phase encoding direction in the antero-posterior direction. Each FOV covered the entire object with a margin of several voxels, to allow SRR of the entire object. As experiments were performed for various aspect factors, we used the minimum number of low-resolution rotations per aspect factor according to Eq. (1). In simulations, Eq. (2) was used to obtain low-resolution data. In all experiments, low-resolution data was reconstructed per slice to high-resolution according to Eq. (7).

3.2. The impact of noise floor on signal accuracy

We characterize the effect of the rectified noise floor on signal accuracy when modulating the aspect factor as well as the strength of diffusion encoding. We use a mono-exponential signal model in water given by

$$S(\alpha, b) = S_0 \cdot \alpha \cdot \exp(-b \cdot D),$$

(15)

where S is the signal in the absence of noise, S_0 is the non-diffusion weighted signal at $\alpha = 1$, b is the strength of the diffusion encoding and D is the diffusion coefficient. To obtain a signal model that includes the contribution of Rician noise, we use the simplified analytical model as described in Eq. (8).

Measurements were performed in a phantom filled with deionized water at room temperature. Data were acquired in a single slice with an in-plane resolution of 1.6×1.6

mm^2 and slice thicknesses between 1.6 mm and 9.6 mm, resulting in aspect factors between 1 and 6. A constant TR of 4 s was used to remove influence of T1-relaxation. A single slice acquisition was used to ensure that the same part of the phantom was imaged for all aspect factors, thereby providing a fair comparison of baseline signals. We used b-values ranging from 0 to 3 ms/μm^2 in steps of 0.3 ms/μm^2 and 5 repetitions. The average measured signal $\bar{S}(\alpha, b)$ was estimated for each aspect factor and b-value in a homogeneous area of the phantom. The non-diffusion weighted signal S_0, diffusivity D and noise level σ were estimated by a least-squares fit of the measured data to Eq. (8) given Eq. (15)

$$\hat{S}_0, \hat{D}, \hat{\sigma} = \underset{S_0, D, \sigma}{arg\ min} \left\| \sqrt{(S_0 \cdot \alpha \cdot \exp(-b \cdot D))^2 + (\sigma \sqrt{\pi/2})^2} - \bar{S}(\alpha, b) \right\|_2^2. \quad (16)$$

The same conditions were reproduced in numerical simulations. The simulations were matched to the observations in the water phantom, and used $S_0 = 1$ and $D = 2.2\ \mu$m^2/ms. Instead of using the analytical model for noise contribution, the noise characteristics were simulated by adding noise sampled from a normal distribution with standard deviation $\sigma = 0.014$. Independent realizations of noise were added to the real and imaginary part of complex signal generated by Eq. (15), whereafter its magnitude was computed. This yielded an SNR of 71 and 426 for $\alpha = 1$ and 6, respectively. The average signal $\bar{S}(\alpha, b)$ was estimated from 10^4 realizations of noise.

Finally, similar simulations were performed within SRR for aspect factors between 1 and 6, deployed on a digital phantom of a cylinder filled with water ($S_0 = 1$ and $D = 2.2\ \mu$m^2/ms). For each aspect factor, a single noisy ($\sigma = 0.014$) set of low-resolution measurements was synthesized using Eq. (2) and reconstructed to high-resolution, after which the average signal $\bar{S}(\alpha, b)$ was computed in a ROI within the cylinder.

The estimate SNFR for each aspect factor, we fitted Eq. (16) to the signal generated by the three methods to obtain σ. From σ and the measured signal $\bar{S}(\alpha, 0)$, we computed SNFR according to Eq. (9). To capture the trend of how SNFR affects the accuracy of signal, we computed the threshold attenuation factor ($\beta_{thr} = b \cdot D$) at which the signal bias exceeds 5%, i.e., $f = \bar{S}(\alpha, b)/S(\alpha, b) > 1.05$, using

$$\beta_{thr} = \log\left(\frac{S_0 \cdot \alpha \cdot \sqrt{f^2 - 1}}{\sigma \sqrt{\pi/2}} \right). \quad (17)$$

The threshold attenuation factors were investigated as a function of aspect factors to visualize how bias thresholds can be extended by imaging at lower resolution.

3.3. Analysis of precision and SNR efficiency

We performed numerical simulations to investigate the relationship between SNR efficiency, aspect factor, T1-relaxation and acquisition time. We expect the SNR efficiency to monotonically increase with the aspect factor when TR is relatively long compared to T1. We used a Shepp-Logan phantom to visualize the effect of regularization on resolution.

Low-resolution data was sampled for aspect factors between 1 and 8 assuming $S \propto \alpha$ and $S_{\alpha=1} = 1$ in the largest part of the phantom, mimicking white matter. Gaussian noise with $\sigma = 0.1$ was added and the SNR efficiency was estimated from the reconstructed images in a homogenous area according to Eq. (11), where we set $\frac{TA_{SRR}}{TA_D} = \frac{N_R}{\alpha}$.

Since the regularization in SRR incurs a blurring, it may also cause an unfair advantage in terms of signal precision. To compensate for this effect, we compared SNR efficiency at similar effective spatial resolutions. A graphical overview of this process is given in Appendix C, Fig. C1. Briefly, we ensure similar effective spatial resolutions by matching the effect of smoothing due to regularization in SRR to gaussian smoothing in direct sampling. We base the matching on a noise-free image that depicts a random texture (each voxel value is drawn form a normal distribution). To obtain directly sampled references, this image was smoothened with 2D Gaussian filters with kernel sizes of 0.40 (weak), 0.47 (moderate) and 0.54 (strong) in units of voxels. The SRR counterparts were created by adjusting the regularization strength (λ' in Eq. (7)) such that the sum of square distances was minimized between the direct and SRR variants. This was done for every combination of kernel size and SRR imaging setup. To avoid the effect of partial SRR sampling at the edges of the FOV, we only matched the region that was covered in all rotations of the FOV. The obtained pairs (smoothing kernel/regularization strength) were used throughout.

The estimation of SNR efficiency was also performed analytically using Eq. (13). As we expect that reducing the TR in SRR leads to reduced SNR efficiency due to T1-relaxation effects, we included T1-weighting in the same equation. We used values of T1 as found in brain at 3T, with T1 = 0.8 s and T1 = 1.6 s to mimic white and gray matter, respectively (Wright et al., 2008). In gray matter, we evaluated SNR efficiency at moderate regularization for TR_D = [5 10 20] s. We set TR_{SRR} to the minimum possible value for each aspect factor (Eq. (14)). However, TR_{SRR} does not necessarily have to assume the minimal value and can be increased at the expense of scan time. Increasing TR_{SRR} will lead to an increased SNR efficiency by promoting T1-relaxation, which is counteracted by a decrease in SNR efficiency as the additional time can be used for averaging the direct acquisition (Eq. (13)). Therefore, we investigated how an increased TR_{SRR} affects the SNR efficiency in both gray and white matter for $\alpha = 8$.

3.3.1.   Analysis of precision in vivo—SNR efficiency was evaluated in a healthy brain in vivo (male, 28 years) for aspect factors between 1 and 6. All experiments used b = [0 0.5] ms/μm^2 with 1 and 10 repetitions respectively, FOV = $220 \times 220 \times 144$ mm^3, TE = 100 ms, partial-Fourier factor = 6/8, an in-plane acceleration factor of 2 (GRAPPA), and bandwidth = 1725 Hz/pixel. Imaging parameters that depended on the aspect factor are summarized in Table 1. All images were reconstructed at a resolution of $1.6 \times 1.6 \times 1.6$ mm^3. To enable a fair comparison with the direct acquisition approach, it was smoothened using 2D Gaussian with a kernel size of 0.47 to yield an effective spatial resolution similar to a moderately regularized SRR. The SNR efficiency was estimated according to Eq. (11) in the central white matter, brainstem, corpus callosum and cerebellar white matter at b = 0.5 ms/μm^2, whereas images acquired at b = 0 ms/μm^2 are shown.

3.4. In vivo dot fraction imaging

We demonstrated the utility of SRR by visualizing the presence of the so-called dot fraction in the brain with a higher contrast than that from direct sampling. The dot fraction is linked to the relative signal that remains at very high b-values when using spherical b-tensor encoding. As anisotropic tissue is effectively attenuated by spherical b-tensor encoding, the remaining signal can be attributed to environments in which the apparent diffusivity is low or zero in all directions (Dhital et al., 2018; Tax et al., 2020).

A healthy volunteer (male, 27 years) was scanned at $b = [0\ 1\ 4]$ ms/μm^2 using 1, 3, and 13 repetitions, $N_R = 8$, resolution $= 1.6 \times 1.6 \times 7.2$ mm^3 ($\alpha = 4.5$), FOV $= 211 \times 211 \times 144$ mm^3, TE $= 120$ ms, TR $= 4200$ ms, TA $= 9{:}31$ min, 2x in-plane acceleration (GRAPPA), partial-Fourier $= 6/8$ and bandwidth $= 1720$ Hz/pixel. A directly sampled set was acquired for comparison with the same settings, at $b = [0\ 1\ 4]$ ms/μm^2 using 1,6, and 33 repetitions, resolution $= 1.6 \times 1.6 \times 1.6$ mm^3, FOV $= 211 \times 211 \times 188$ mm^3, TR $= 14{,}200$ ms and TA $= 9{:}28$ min. As regularized SRR reduces the spatial resolution slightly, another directly sampled set was acquired with the same settings at a resolution of $1.8 \times 1.8 \times 1.8$ mm^3 with FOV $= 238 \times 238 \times 162$ mm^3.

The TR was set to the minimum value possible without the use of through-plane acceleration. To reduce the impact of system heating and drift, we used simultaneous multislice acceleration of factor 2 and randomized the b-values across volumes (Hutter et al., 2018; Szczepankiewicz et al., 2021; Vos et al., 2017). All raw data were denoised using Marchenko-Pastur principle component analysis (Cordero-Grande et al., 2019; Tournier et al., 2019). The low-resolution images were reconstructed at a resolution of $1.6 \times 1.6 \times 1.6$ mm^3 using SRR. A moderate regularization of $\lambda' = 0.026$ was used.

We assume a dot compartment with signal fraction f_{dot} and isotropic diffusivity D_{dot} equal to zero, accompanied by a fraction of other tissue $(1 - f_{\text{dot}})$ with non-zero isotropic diffusivity. Assuming Gaussian diffusion and no exchange, the diffusion-weighted signal $S(b)$ is given by (Dhital et al., 2018; Tax et al., 2020)

$$S(b) = S_0 \cdot (f_{\text{dot}} \exp(-b \cdot D_{\text{dot}}) + (1 - f_{\text{dot}}) \exp(-b \cdot D_{\text{other}})), \tag{18}$$

where S_0 is the non-diffusion weighted signal. For very high b-values, the $\exp(-b_{\text{high}} \cdot D_{\text{other}}) \approx 0$, and only signal in the dot compartment remains since $D_{\text{dot}} \approx 0$, which simplifies Eq. (18) and allows us to rewrite to

$$S(b) \approx S_0 \cdot f_{\text{dot}} \text{ and } f_{\text{dot}} \approx \frac{S(b)}{S_0}. \tag{19}$$

The resulting signal map is weighted by both S_0 (T2-weighting) and f_{dot} (influenced by density of cells exhibiting restricted diffusion in all directions). Note that, under these assumptions, f_{dot} only gives an upper limit on the true value as other tissues as well as the rectified noise floor may contribute to the remaining signal.

Both the super-resolved and directly acquired high-resolution images were averaged for $b = 4$ ms/μm^2. The gray-to-white matter signal ratio was calculated between gray matter voxels in the cerebellar cortex and white matter voxels in the cerebellum. f_{dot} was calculated using Eq. (19). We compared the images to a corresponding T1-weighted morphological scan and to a similar Nissl-stained contrast that emphasize neurons found in brain histology from the BigBrain atlas (Amunts et al., 2013).

4. Results

4.1. The impact of noise floor on signal accuracy

Fig. 2 shows that the rectified noise floor causes an overestimation of the signal at high b-values where SNR is low (Fig. 2a). The bias can be partially avoided by using higher aspect factors to increase the SNFR. This means that higher b-values can be employed before reaching a given signal bias. In the experiments, the measured signal at $\alpha = 1$ exhibited an SNFR of 60, which could be increased to 360 at $\alpha = 6$. By doing so, the maximal b-value could be increased from 1.32 to 2.13 ms/μm^2 before reaching the 5% signal bias in water at room temperature (Fig. 2b). More generally, increasing the aspect factor from 1 to 6 enabled the attenuation coefficient to become 62% larger (Eq. (17)). Both single voxel and SRR simulations agreed with measurements. Note that TR was here constant across measurements to avoid contributions from T1-weighting.

4.2. Analysis of precision and SNR efficiency

Fig. 3 shows the SNR efficiency for SRR with different aspect factors for the three levels of regularization (see Appendix C, Fig. C2 for the regularization values). SNR efficiency generally increases with aspect factor. For example, the SNR efficiency is 1.6 for an aspect factor 8 when using moderate regularization. As expected, strong regularization leads to higher SNR efficiency, but may cause reduced spatial resolution and blurred edges. Conversely, weak regularization may substantially inflate the effects of noise, resulting in SNR efficiency below 1 for all aspect factors. As analytical and numerical results agree, the analytical expression can guide the regularization accurately. Note that this analysis does not include T1-relaxation effects.

Fig. 4 shows the effect of T1-relaxation on SNR efficiency, which decreases with TR_D (Fig. 4a). This effect is more evident for higher aspect factors. For example, for aspect factor 8, SNR efficiency reduces from 1.6 to 0.6 as we go from the case where T1-relaxation is neglected ($TR_D \gg T1$) to a case where $TR_D = 5$ s for a T1 value mimicking gray matter. SRR in this case leads to a reduced SNR efficiency, rather than an increase. However, when $TR_D \gg T1$, T1-effects are less evident and SNR efficiency is still above unity for SRR. This situation occurs for example when TR_D is relatively high and/or low aspect factors are used. Note that TR_{SRR} can be increased to promote T1-relaxation. However, this leads to an increased scan time, which benefits the direct acquisition as this time can be used for signal averaging. As illustrated for an aspect factor of 8, we see it is beneficial to set TR_{SRR} to a minimum when $TR_{SRR}/T1 > 1.25$ (Fig. 4b). Until this point is reached, the benefits of rapid signal-sampling outweigh the benefits of prolonged T1-relaxation in terms of SNR efficiency.

4.2.1.   Analysis of precision in vivo—Fig. 5 shows the results of SRR in vivo for different aspect factors at similar effective resolution. It shows that SRR in vivo is feasible, as resolution is regained for all aspect factors. Ringing artefacts are observed near high-contrast transitions, such as around the ventricles (Gottlieb et al., 1992). As the aspect factor increases and the TR_{SRR} is reduced, T1-weighting becomes more visible. This can for example be observed in the contrast decrease between white matter and CSF, with T1 of 0.8 s and 4.5 s, respectively (Shin et al., 2009).

Fig. 6 shows SNR efficiency for measured and simulated data in white matter at $b = 0.5$ ms/μm^2. In central white matter and the corpus callosum, experiments and simulations (T1 = 0.8 s) follow similar trends of increased SNR efficiency with aspect factor, up to $\alpha = 5$. In the brainstem and cerebellar white matter, simulations overestimate the SNR efficiency for all aspect factors and the trends are more similar to simulations with a longer T1 of 1.6 s, mimicking gray matter. A potential explanation for this trend is that ROIs in both brainstem and cerebellum contain gray matter with relatively long T1, thereby reducing SNR efficiency according to Eq. (13). Regardless, the SNR efficiency is still above unity for all cases, meaning that SRR has a higher precision than its directly sampled equivalent for this protocol.

4.3.   In vivo dot fraction imaging

Fig. 7 shows results for the dot fraction imaging. Direct sampling at 1.6 mm^3 isotropic leads to poor image contrast throughout the brain, as the image is dominated by the noise floor. When sampling directly at a resolution of 1.8 mm^3 isotropic, the cerebellar cortex becomes faintly visible, but the contrast is still poor. On the other hand, SRR enables a vastly improved image contrast where the cortex of the cerebrum is visible with prominent signal. The contrast ratio between cortex and white matter of the cerebellum are 1.86 for SRR versus 1.06 and 1.31 for direct sampling at 1.6 mm^3 and 1.8 mm^3 isotropic, respectively. Fig. 8 shows that a part of this contrast is due to T2 effects, as the map of f_{dot} shows a less pronounced difference between cerebellar white- and gray matter. A similar contrast is observed in neuron-stained brain histology, where the cerebellar cortex shows high density of neurons.

5.   Discussion

In this work, we investigated the performance of super-resolution reconstruction and its impact on signal accuracy and precision in diffusion MRI. We found that SRR improved accuracy in a challenging application that would not be feasible without it (Fig. 7). We also expect that the increase in signal accuracy by SRR will lead to an improved accuracy in parameters estimated from dMRI data (Jeurissen et al., 2018).

SRR improves the accuracy of the diffusion-weighted signal by avoiding sampling signal that is biased by the rectified noise floor. This gain is linear with voxel volume in terms of SNFR, which allows the use of stronger diffusion encoding before a signal bias is introduced (Fig. 2). An alternative method to improve accuracy is averaging complex data, which is challenging due to the need of phase correction (Eichner et al., 2015; Prah et al., 2010). The presented methodology uses the magnitude signal, which allows us to avoid this challenge.

Although SRR and complex averaging are, in principle, compatible, the investigation of such an approach was outside the scope of this work. Moreover, as signal bias is mostly avoided rather than corrected, SRR does not have to rely on prior knowledge of the data distribution. This differs from postprocessing methods that have been proposed to remove the bias (Andersson, 2008; Kristoffersen, 2007; Veraart et al., 2013, 2011) and can be challenging due to, for example, non-trivial coil-combination methods. An alternative SRR technique is g-slider (Setsompop et al., 2018) where signal is acquired in slabs with different excitation profiles. Although this approach has excellent SNR efficiency, SRR based on rotating FOVs can be deployed with off-the-shelf sequences, whereas g-slider requires a specialized pulse sequence. Nevertheless, a combination of g-slider, spherical b-tensor encoding and high spatial resolution may enable even better resolution, contrast and/or shorter acquisition times.

We found that SNR efficiency can be preserved or even improved using SRR, depending on the aspect factor, T1-weighting, and regularization. The established analytical formulation for SNR efficiency (Eq. (13)) may help setting the regularization strength to decide the tradeoff between SNR efficiency and effective spatial resolution (Fig. 3). Simulations show that SNR efficiency is generally increased for large aspect factors, however, effects of T1-relaxation become dominant at sufficiently high aspect factors (Fig. 4). Previous studies suggested larger aspect factor would lead to larger improvements in SNR efficiency (Van Steenkiste et al., 2016). However, this was under the assumption of complete T1-recovery. The effects of T1-relaxation counteract the SRR improvement when relatively short repetition times are used, which has become more common since the simultaneous multi-slice techniques were made available (Barth et al., 2016). Furthermore, the use of thicker slices increases sensitivity to motion-induced gradient imbalances leading to intravoxel dephasing and therefore loss of signal, which may partially explain the mismatch between simulations and in vivo results (Fig. 6) and put a practical limit on the maximum achievable aspect factor (Gumus et al., 2014). Generally, we see that the potential precision increase is heavily dependent on acquisition, reconstruction, and intrinsic tissue parameters.

The gain in precision from SRR was only modest, while accuracy can be substantially improved by SRR in low SNR scenarios. In practice, this leads to an improved contrast, as the use of larger voxels will reduces the influence from the noise floor and therefore allow for a better differentiation of true signal differences. We demonstrated this in vivo for so called dot-fraction imaging, which visualizes densely packed cells located in the cerebellar cortex and was first demonstrated on a custom MRI system with ultra-strong gradients and at a poor spatial resolution (Tax et al., 2020). Using SRR, we enabled high-resolution imaging at 1.6 mm^3 isotropic at $b = 4$ ms/μm^2 with spherical tensor encoding, where the elevated SNFR causes contrast to be substantially improved compared to direct sampling (Fig. 7). If a linear relationship with voxel volume is assumed, the same SNFR is expected for direct sampling only at 2.7 mm^3 isotropic. We believe that this approach facilitates high-resolution studies of neurodegenerative diseases where isotropic tissues are affected. Current acquisition times are just below 10 min, and the sampling scheme can be further optimized to comply with clinical routine.

Under high-SNR conditions, the simplicity of direct sampling may outweigh the benefits of SRR. However, a remaining benefit of SRR is that it may be used for acceleration when used in combination with a diffusion model (Jeurissen et al., 2018; Van Steenkiste et al., 2016). Acceleration can be achieved by subsampling the diffusion encoding directions or b-tensors over the acquired low-resolution images, such that directional information needed to fit the model can be acquired in a shorter time.

We have identified five primary limitations of the present study. First, to compare regularized SRR and direct sampling at similar effective resolution, we matched the effect of regularization strength and Gaussian smoothing kernel size. However, matching spatial resolutions across filtering and/or reconstruction methods remains a challenge as metrics of spatial resolution are lacking (den Dekker and van den Bos, 1997). Notably, both regularization and smoothing always lead to an effective resolution that is lower than the original or nominal resolution. As such, sampling directly at the effective resolution would increase both SNFR (Fig. 2) and SNR efficiency (Scouten et al., 2006) compared to sampling at the original resolution and performing post-hoc smoothing. For an evaluation of effective spatial resolution and therefore better assessment of SRR, suitable phantoms could be used in future work (Fellner et al., 2001; Price et al., 1990). Second, our analysis does not include motion and eddy current correction in the SRR model, while perfect registration in SRR is of importance to obtain non-blurry high-resolution results. Including such corrections could improve the sharpness of in vivo results. Third, we used the identity matrix as the regularization matrix (Eq. (7)). Using data-driven regularization that depends on image content can have benefits, such as edge-preservation (Khattab et al., 2020). Forth, the multi-stack rotation SRR model suffers from signal over- and undershoot problems at contrast-rich regions due to regularization, as more thoroughly described by (Shilling et al., 2009). This could be a confounder in quantitative measures extracted from the images, for example when estimating the dot-fraction. Fifth, while concomitant gradients are compensated in the diffusion encoding (Szczepankiewicz et al., 2019b), the read-out gradients are rotated along with the FOV and their concomitant gradients may cause rotation-dependent artefacts (Du et al., 2002).

In conclusion, we have presented a comprehensive analysis of SRR that outlined the major features influencing the precision and accuracy of the diffusion-weighted signal. The primary benefit of SRR is the improved signal-to-noise-floor ratio, but precision can also be improved given beneficial imaging conditions. We showcased the use of SRR in a challenging combination of high resolution and spherical tensor encoding with ultra-high b-values, where SRR can suppress noise floor effects and recover high signal accuracy. We share the tools developed herein in open-source and expect that the present work will support future experimental design, such that both acquisition and reconstruction parameters can be optimized for specialized purposes.

Acknowledgments

Appendix A. Derivation of SNR efficiency

We derive ρ as given in Eq. (13). We start from Eq. (11), stating

$$\rho = \frac{SNR_{SRR}}{SNR_D} \cdot \sqrt{\frac{TA_D}{TA_{SRR}}}. \tag{A1}$$

To find our final expression, we express both SNR_{SRR} and SNR_D in terms of SNR_{LR}, the SNR of a low-resolution acquisition. For all cases,

$$\overline{S}_x \propto V_x \cdot (1 - \exp(-TR_x/T1)), \tag{A2}$$

where $\overline{S}$ is the mean signal, V is the voxel size and x the reflecting component, i.e. LR, SRR or D (direct sampling). As $V_{LR} = \alpha \cdot V_D$, we rewrite Eq. (A2) to

$$\overline{S}_D = \frac{\overline{S}_{LR}}{\alpha} \cdot \frac{1 - \exp(-TR_D/T1)}{1 - \exp(-TR_{LR}/T1)}. \tag{A3}$$

As σ is independent of voxel size, $\sigma_D = \sigma_{LR}$ (Edelstein et al., 1986), and combining both Eq. (A3) and Eq. (10) gives

$$SNR_D = \frac{1}{\alpha} \cdot \frac{1 - \exp(-TR_D/T1)}{1 - \exp(-TR_{LR}/T1)} \cdot SNR_{LR}. \tag{A4}$$

Along the same way we can find an expression for SNR_{SRR}. As $TR_{SRR} = TR_{LR}$, T1-effects are the same for both measurements, and we use Eq. (A2) to see that

$$\overline{S}_{SRR} = \frac{\overline{S}_{LR}}{\alpha}. \tag{A5}$$

We use the definition of the noise propagation factor κ of Eq. (12) and Eq. (10) to find

$$SNR_{SRR} = \frac{1}{\alpha \cdot \kappa} \cdot SNR_{LR}. \tag{A6}$$

As for SRR the repetition time decreases, but the number of slice orientations increases with a factor N_R, we find

$$\frac{TA_{SRR}}{TA_D} = \frac{N_R \cdot TR_{SRR}}{TR_D}. \tag{A7}$$

Combing Eq. (A4), Eq. (A6) and Eq. (A7) in Eq. (A1), gives the SNR efficiency ρ defined in Eq. (13), being

$$\rho = \frac{1}{\kappa} \cdot \frac{1 - \exp(-\mathrm{TR_{SRR}}/T1)}{1 - \exp(-\mathrm{TR_D}/T1)} \cdot \sqrt{\frac{\mathrm{TR_D}}{N_R \cdot \mathrm{TR_{SRR}}}} . \tag{A8}$$

Appendix B. Gradient waveforms

Conventional, or linear diffusion encoding yields a pair of trapezoidal pulsed field gradient on each side of the refocusing pulse in a spin-echo sequence (Stejskal and Tanner, 1965). In this work, we used more advanced, spherical encoding where all three gradients are continuously used (Szczepankiewicz et al., 2019a). The gradient waveforms we used are depicted in Fig. B1.
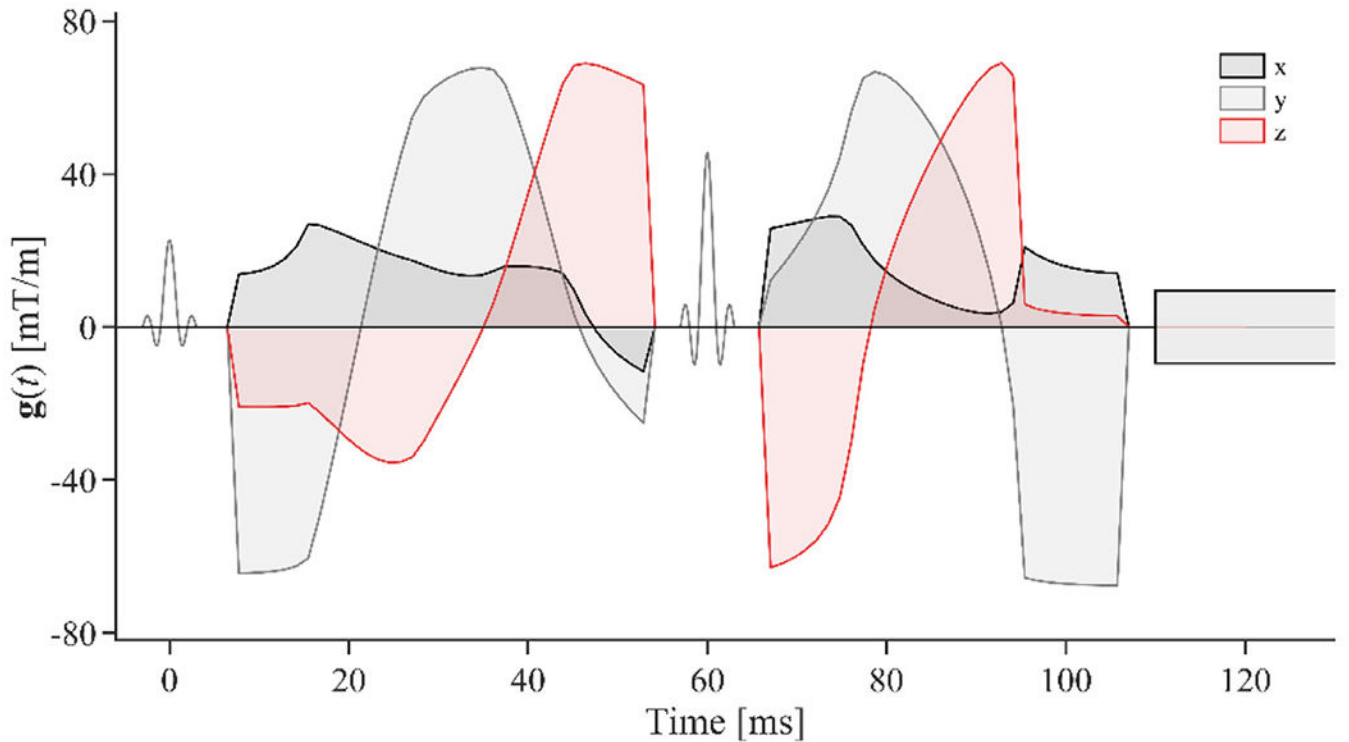
Fig. B1.
Optimized gradient waveforms were used to yield spherical tensor encoding (STE) in all practical experiments, and the one used for in vivo measurements is shown under conditions that yield $b = 4.0$ ms/μm^2. Radiofrequency pulses (RF) are depicted in between gradients. This waveform and other resources related to the free waveforms sequence are available at https://github.com/filip-szczepankiewicz/fwf_seq_resources.

Appendix C. Matching the effective spatial resolution

The effective spatial resolution between super-resolution reconstructed and directly sampled images are matched, as illustrated in Fig. C1. The regularization values used for different aspect factors are given in Fig. C2.
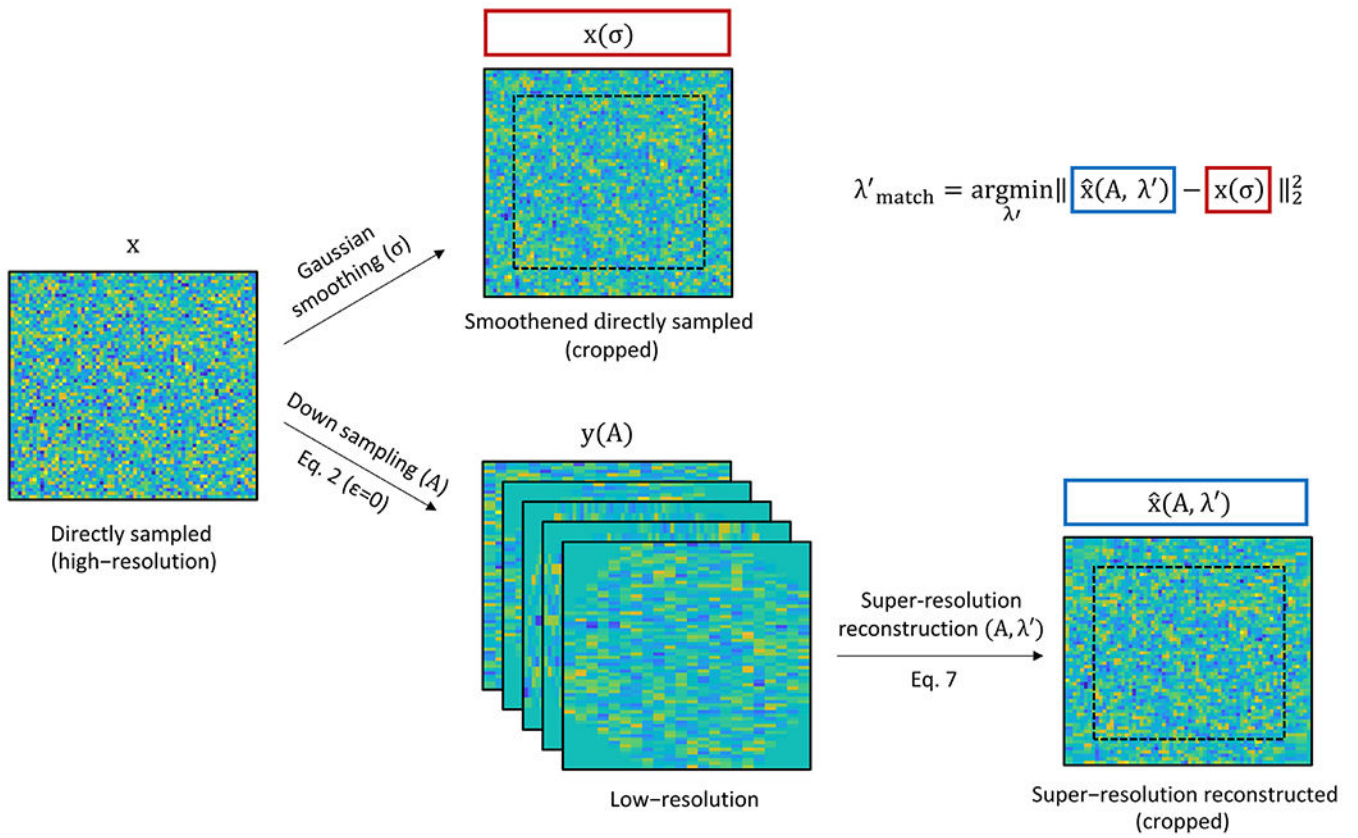
$$\lambda'_{\text{match}} = \underset{\lambda'}{\arg\min} \| \hat{x}(A, \lambda') - x(\sigma) \|^2_2$$

Fig. C1.

Matching the effective spatial resolution of directly sampled and super-resolution reconstructed images. A noise-free high-resolution image x is smoothened with a given Gaussian kernel size σ and separately down sampled to low-resolution y with a given down sampling matrix A (A is dependent on the aspect factor). The super-resolution reconstructed image $\hat{x}(A, \lambda')$ has similar effective spatial resolution as the smoothened directly sampled image x(σ) for a regularization strength λ' that minimizes the sum of squares distance between them. Images are cropped before comparison to avoid influence of edge-artefacts.
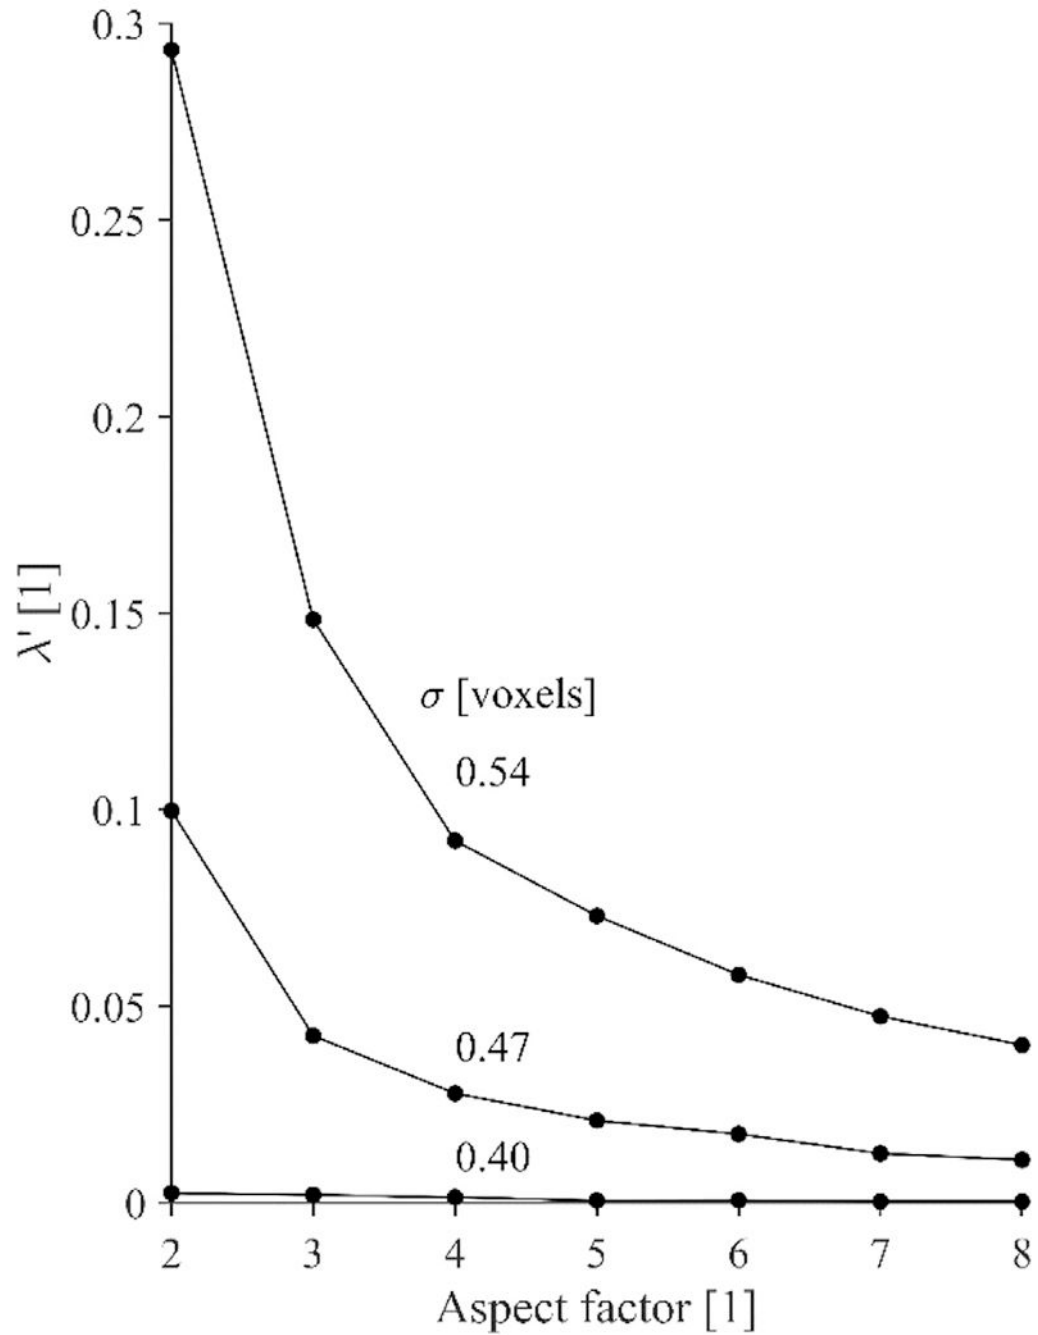
Fig. C2.
Values of regularization strength λ' matched with Gaussian kernel sizes σ in terms of induced effective spatial resolution for different aspect factors. λ' decreases with increasing aspect factor, with a smaller decrease for higher aspect factors. λ' increases with increasing σ.

Data availability statement

The data that support the findings of this study are not available due to ethical restrictions.

The super-resolution reconstruction framework is available in open source at https://github.com/filip-szczepankiewicz/Vis_NIMG_2021/.

References

Amunts K, Lepage C, Borgeat L, Mohlberg H, Dickscheid T, Rousseau MÉ, Bludau S, Bazin PL, Lewis LB, Oros-Peusquens AM, Shah NJ, Lippert T, Zilles K, Evans AC, 2013. BigBrain: an ultrahigh-resolution 3D human brain model. Science 340 (80), 1472–1475. doi:10.1126/science.1235381. [PubMed: 23788795]

Andersson JLR, 2008. Maximum a posteriori estimation of diffusion tensor parameters using a Rician noise model: why, how and but. Neuroimage 42, 1340–1356. doi:10.1016/j.neuroimage.2008.05.053. [PubMed: 18602480]

Barth M, Breuer F, Koopmans PJ, Norris DG, Poser BA, 2016. Simultaneous multislice (SMS) imaging techniques. Magn. Reson. Med 75, 63–81. doi:10.1002/mrm.25897. [PubMed: 26308571]

BS ISO 5725-1, 1994. Accuracy (trueness and precision) of measurement methods and results - Part 1: general principles and definitions. p. 1.

Cárdenas-Blanco A, Tejos C, Irarrazaval P, Cameron I, 2008. Noise in magnitude magnetic resonance images. Concepts Magn. Reson. Part A 32A, 409–416. doi:10.1002/cmr.a.20124.

Cordero-Grande L, Christiaens D, Hutter J, Price AN, Hajnal JV, 2019. Complex diffusion-weighted image estimation via matrix recovery under general noise models. Neuroimage 200, 391–404. doi:10.1016/j.neuroimage.2019.06.039. [PubMed: 31226495]

den Dekker AJ, Sijbers J, 2014. Data distributions in magnetic resonance images: a review. Phys. Medica 30, 725–741. doi:10.1016/j.ejmp.2014.05.002.

den Dekker AJ, van den Bos A, 1997. Resolution: a survey. J. Opt. Soc. Am. A 14, 547–557. doi:10.1364/JOSAA.14.000547.

Dhital B, Kellner E, Kiselev VG, Reisert M, 2018. The absence of restricted water pool in brain white matter. Neuroimage 182, 398–406. doi:10.1016/j.neuroimage.2017.10.051. [PubMed: 29129672]

Du YP, Joe Zhou X, Bernstein MA, 2002. Correction of concomitant magnetic field-induced image artifacts in nonaxial echo-planar imaging. Magn. Reson. Med 48, 509–515. doi:10.1002/mrm.10249. [PubMed: 12210916]

Edelstein WA, Glover GH, Hardy CJ, Redington RW, 1986. The intrinsic signal-to-noise ratio in NMR imaging. Magn. Reson. Med 3, 604–618. doi:10.1002/mrm.1910030413. [PubMed: 3747821]

Eichner C, Cauley SF, Cohen-Adad J, Möller HE, Turner R, Setsompop K, Wald LL, 2015. Real diffusion-weighted MRI enabling true signal averaging and increased diffusion contrast. Neuroimage 122, 373–384. doi:10.1016/j.neuroimage.2015.07.074. [PubMed: 26241680]

Elad M, Feuer A, 1997. Restoration of a single superresolution image from several blurred, noisy, and undersampled measured images. IEEE Trans. Image Process 6, 1646–1658. doi:10.1109/83.650118. [PubMed: 18285235]

Fellner C, Müller W, Georgi J, Taubenreuther U, Fellner FA, Kalender WA, 2001. A high-resolution phantom for MRI. Magn. Reson. Imaging 19, 899–904. doi:10.1016/S0730-725X(01)00397-6. [PubMed: 11551732]

Fukutani Y, Cairns NJ, Rossor MN, Lantos PL, 1996. Purkinje cell loss and astrocytosis in the cerebellum in familial and sporadic Alzheimer's disease. Neurosci. Lett 214, 33–36. doi:10.1016/0304-3940(96)12875-5. [PubMed: 8873125]

Gottlieb D, Shu C-W, Solomonoff A, Vandeven H, 1992. On the Gibbs phenomenon I: recovering exponential accuracy from the Fourier partial sum of a nonperiodic analytic function. J. Comput. Appl. Math 43, 81–98.

Greenspan H, 2009. Super-resolution in medical imaging. Comput. J 52, 43–63. doi:10.1093/comjnl/bxm075.

Greenspan H, Oz G, Kiryati N, Peled S, 2002. MRI inter-slice reconstruction using super-resolution. Magn. Reson. Imaging 20, 437–446. doi:10.1016/S0730-725X(02)00511-8. [PubMed: 12206870]

Gudbjartsson H, Patz S, 1995. The Rician distribution of noisy MRI data. Magn. Reson. Med 34, 910–914. doi:10.1002/mrm.1910340618. [PubMed: 8598820]

Gumus K, Keating B, Poser BA, Armstrong B, Chang L, Maclaren J, Prieto T, Speck O, Zaitsev M, Ernst T, 2014. Prevention of motion-induced signal loss in diffusion-weighted echo-planar imaging by dynamic restoration of gradient moments. Magn. Reson. Med 71. doi:10.1002/mrm.24857, 2006–2013.

Hutter J, Nilsson M, Christiaens D, Schneider T, Price AN, Hajnal JV, Szczepankiewicz F, 2018. Highly efficient diffusion MRI by slice-interleaved free-waveform imaging (SIFI). In: Proceedings of the 26th Annual Meeting of ISMRM. Paris, France.

Jensen JH, Helpern JA, Ramani A, Lu H, Kaczynski K, 2005. Diffusional kurtosis imaging: the quantification of non-gaussian water diffusion by means of magnetic resonance imaging. Magn. Reson. Med 53, 1432–1440. doi:10.1002/mrm.20508. [PubMed: 15906300]

Jeurissen B, Westin C-F, Sijbers J, Szczepankiewicz F, 2018. Improved precision and accuracy in q-space trajectory imaging by model-based superresolution reconstruction. Abstr. Submitt to ISMRM 2018.

Jezzard P, Barnett AS, Pierpaoli C, 1998. Characterization of and correction for eddy current artifacts in echo planar diffusion imaging. Magn. Reson. Med 39, 801–812. doi:10.1002/mrm.1910390518. [PubMed: 9581612]

Jones DK, Basser PJ, 2004. Squashing peanuts and smashing pumpkins": how noise distorts diffusion-weighted MR data. Magn. Reson. Med 52, 979–993. doi:10.1002/mrm.20283. [PubMed: 15508154]

Khattab MM, Zeki AM, Alwan AA, Badawy AS, 2020. Regularization-based multi-frame super-resolution: a systematic review. J. King Saud Univ. - Comput. Inf. Sci 32, 755–762. doi:10.1016/j.jksuci.2018.11.010.

Kristoffersen A, 2007. Optimal estimation of the diffusion coefficient from non-averaged and averaged noisy magnitude data. J. Magn. Reson 187, 293–305. doi:10.1016/j.jmr.2007.05.004. [PubMed: 17572124]

Lätt J, Nilsson M, Malmborg C, Rosquist H, Wirestam R, Ståhlberg F, Topgaard D, Brockstedt S, 2007. Accuracy of q-space related parameters in MRI: simulations and phantom measurements. IEEE Trans. Med. Imaging 26, 1437–1447. doi:10.1109/TMI.2007.907278. [PubMed: 18041259]

Lundell H, Nilsson M, Dyrby TB, Parker GJM, Cristinacce PLH, Zhou F-L, Topgaard D, Lasič S, 2019. Multidimensional diffusion MRI with spectrally modulated gradients reveals unprecedented microstructural detail. Sci. Rep 9, 9026. doi:10.1038/s41598-019-45235-7. [PubMed: 31227745]

Plenge E, Poot DHJ, Bernsen M, Kotek G, Houston G, Wielopolski P, Van Der Weerd L, Niessen WJ, Meijering E, 2012. Super-resolution methods in MRI: can they improve the trade-off between resolution, signal-to-noise ratio, and acquisition time? Magn. Reson. Med 68. doi:10.1002/mrm.24187, 1983–1993.

Poot DHJ, Jeurissen B, Bastiaensen Y, Veraart J, Van Hecke W, Parizel PM, Sijbers J, 2013. Super-resolution for multislice diffusion tensor imaging. Magn. Reson. Med 69, 103–113. doi:10.1002/mrm.24233. [PubMed: 22411778]

Prah DE, Paulson ES, Nencka AS, Schmainda KM, 2010. A simple method for rectified noise floor suppression: phase-corrected real data reconstruction with application to diffusion-weighted imaging. Magn. Reson. Med 64, 418–429. doi:10.1002/mrm.22407. [PubMed: 20665786]

Price RR, Axel L, Morgan T, Newman R, Perman W, Schneiders N, Selikson M, Wood M, Thomas SR, 1990. Quality assurance methods and phantoms for magnetic resonance imaging: report of AAPM nuclear magnetic resonance Task Group No. 1. Med. Phys 17, 287–295. doi:10.1118/1.596566. [PubMed: 2333055]

Scherrer B, Gholipour A, Warfield SK, 2012. Super-resolution reconstruction of diffusion-weighted images from distortion compensated orthogonal anisotropic acquisitions. Proc. Work. Math. Methods Biomed. Image Anal 2012, 249–254. doi:10.1109/MMBIA.2012.6164759.

Scouten A, Papademetris X, Constable RT, 2006. Spatial resolution, signal-to-noise ratio, and smoothing in multi-subject functional MRI studies. Neuroimage 30, 787–793. doi:10.1016/j.neuroimage.2005.10.022. [PubMed: 16343951]

Setsompop K, Fan Q, Stockmann J, Bilgic B, Huang S, Cauley SF, Nummenmaa A, Wang F, Rathi Y, Witzel T, Wald LL, 2018. High-resolution in vivo diffusion imaging of the human brain with

generalized slice dithered enhanced resolution: simultaneous multislice (gSlider-SMS). Magn. Reson. Med 79, 141–151. doi:10.1002/mrm.26653. [PubMed: 28261904]

Shilling RZ, Robbie TQ, Bailloeul T, Mewes K, Mersereau RM, Brummer ME, 2009. A super-resolution framework for 3-D high-resolution and high-contrast imaging using 2-D multislice MRI. IEEE Trans. Med. Imaging 28, 633–644. doi:10.1109/TMI.2008.2007348. [PubMed: 19272995]

Shin W, Gu H, Yang Y, 2009. Fast high-resolution T1 mapping using i nversion-recovery look-locker echo-planar imaging at steady state: optimization for accuracy and reliability. Magn. Reson. Med 61, 899–906. doi:10.1002/mrm.21836. [PubMed: 19195021]

Sjölund J, Szczepankiewicz F, Nilsson M, Topgaard D, Westin CF, Knutsson H, 2015. Constrained optimization of gradient waveforms for generalized diffusion encoding. J. Magn. Reson 261, 157–168. doi:10.1016/j.jmr.2015.10.012. [PubMed: 26583528]

Souza A, Senn R, 2008. Model-based super-resolution for MRI. In: 2008 30th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, pp. 430–434. doi:10.1109/IEMBS.2008.4649182.

Stejskal EO, Tanner JE, 1965. Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient. J. Chem. Phys 42, 288–292. doi:10.1063/1.1695690.

Szczepankiewicz F, Sjölund J, Ståhlberg F, Lätt J, Nilsson M, 2019a. Tensor-valued diffusion encoding for diffusional variance decomposition (DIVIDE): technical feasibility in clinical MRI systems. PLoS ONE 14, e0214238. [PubMed: 30921381]

Szczepankiewicz F, Westin C-F, Nilsson M, 2021. Gradient waveform design for tensor-valued encoding in diffusion MRI. J. Neurosci. Methods 348, 109007. doi:10.1016/j.jneumeth.2020.109007 [PubMed: 33242529]

Szczepankiewicz F, Westin C-F, Nilsson M, 2019b. Maxwell-compensated design of asymmetric gradient waveforms for tensor-valued diffusion encoding. Magn. Reson. Med 82, 1424–1437. doi:10.1002/mrm.27828. [PubMed: 31148245]

Tax CMW, Szczepankiewicz F, Nilsson M, Jones DK, 2020. The dot-compartment revealed? Diffusion MRI with ultra-strong gradients and spherical tensor encoding in the living human brain. Neuroimage 210, 116534. doi:10.1016/j.neuroimage.2020.116534. [PubMed: 31931157]

Tikhonov AN, Arsenin VY, 1977. Solutions of ill-posed problems.

Tournier J-D, Smith R, Raffelt D, Tabbara R, Dhollander T, Pietsch M, Christiaens D, Jeurissen B, Yeh C-H, Connelly A, 2019. MRtrix3: a fast, flexible and open software framework for medical image processing and visualisation. Neuroimage 202, 116137. doi:10.1016/j.neuroimage.2019.116137. [PubMed: 31473352]

Van Reeth E, Tham I, Tan CH, Poh CL, 2012. Super-resolution in magnetic resonance imaging: a review. Concepts Magn. Reson. Part A 40A, 306–325. doi:10.1002/cmr.a.21249.

Van Steenkiste G, Jeurissen B, Veraart J, Den Dekker AJ, Parizel PM, Poot DHJ, Sijbers J, 2016. Super-resolution reconstruction of diffusion parameters from diffusion-weighted images with different slice orientations. Magn. Reson. Med 75, 181–195. doi:10.1002/mrm.25597. [PubMed: 25613283]

Veraart J, Rajan J, Peeters RR, Leemans A, Sunaert S, Sijbers J, 2013. Comprehensive framework for accurate diffusion MRI parameter estimation. Magn. Reson. Med 70, 972–984. doi:10.1002/mrm.24529. [PubMed: 23132517]

Veraart J, Van Hecke W, Sijbers J, 2011. Constrained maximum likelihood estimation of the diffusion kurtosis tensor using a Rician noise model. Magn. Reson. Med 66, 678–686. doi:10.1002/mrm.22835. [PubMed: 21416503]

Vos SB, Tax CMW, Luijten PR, Ourselin S, Leemans A, Froeling M, 2017. The importance of correcting for signal drift in diffusion MRI. Magn. Reson. Med 77, 285–299. doi:10.1002/mrm.26124. [PubMed: 26822700]

Westin C-F, Knutsson H, Pasternak O, Szczepankiewicz F, Özarslan E, van Westen D, Mattisson C, Bogren M, O'Donnell LJ, Kubicki M, Topgaard D, Nilsson M, 2016. Q-space trajectory imaging for multidimensional diffusion MRI of the human brain. Neuroimage 135, 345–362. doi:10.1016/j.neuroimage.2016.02.039. [PubMed: 26923372]

Wright PJ, Mougin OE, Totman JJ, Peters AM, Brookes MJ, Coxon R, Morris PE, Clemence M, Francis ST, Bowtell RW, Gowland PA, 2008. Water proton T1 measurements in brain tissue at 7, 3, and 1.5 T using IR-EPI, IR-TSE, and MPRAGE: results and optimization. MAGMA 21, 121–130. doi:10.1007/s10334-008-0104-8. [PubMed: 18259791]

Xia G, McFarland KN, Wang K, Sarkar PS, Yachnis AT, Ashizawa T, 2013. Purkinje cell loss is the major brain pathology of spinocerebellar ataxia type 10. J. Neurol. Neurosurg. Psychiatry doi:10.1136/jnnp-2013-305080.

Yang Z, He P, Zhou J, Wu X, 2018. Non-local diffusion-weighted image super-resolution using collaborative joint information. Exp. Ther. Med 15, 217–225. doi:10.3892/etm.2017.5430. [PubMed: 29387188]
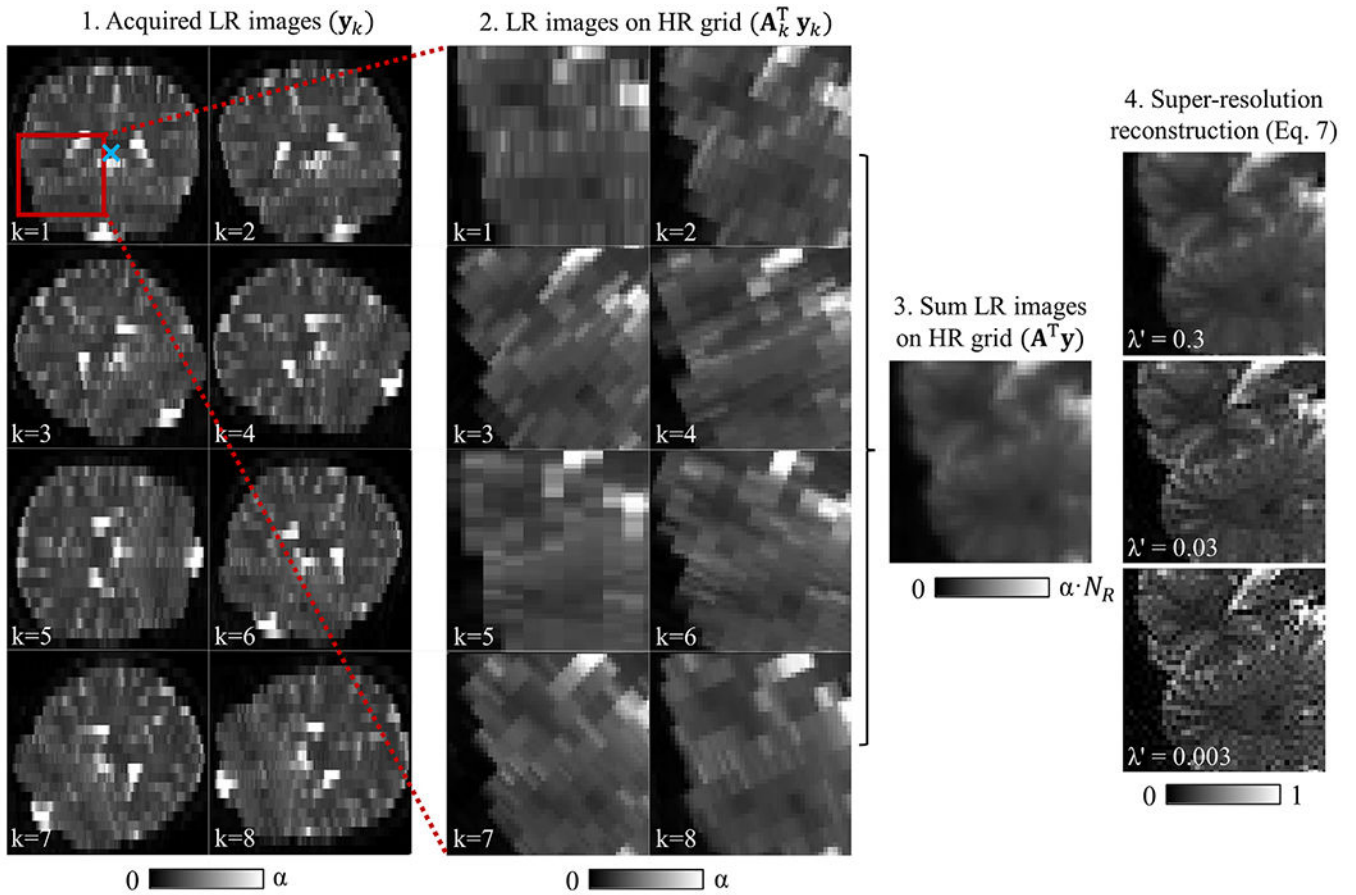
Fig. 1.
Illustration of the super-resolution reconstruction process in coronal view. Step 1: Multiple images are acquired with low through-plane resolution rotated around the anterior-posterior axis indicated by the blue cross. Step 2: The images are up-sampled to the high-resolution grid by application of the up-sampling operator to each individual image. Step 3: The joint up-sampling operator results in the sum of the individual images, i.e., a smooth image on the high-resolution grid. Step 4: The sharpening operator is applied to finalize the reconstruction process and obtain a high-resolution image. The regularization parameter λ' determines the trade-off between resolution and noise propagation; a higher λ' leads to a blurrier, but less noisy image. Steps 2 to 4 are shown as a magnified view of the region indicated by the red square. Color bars show how the intensity scale changes with each step. α is the aspect factor and N_R is the number of low-resolution images.
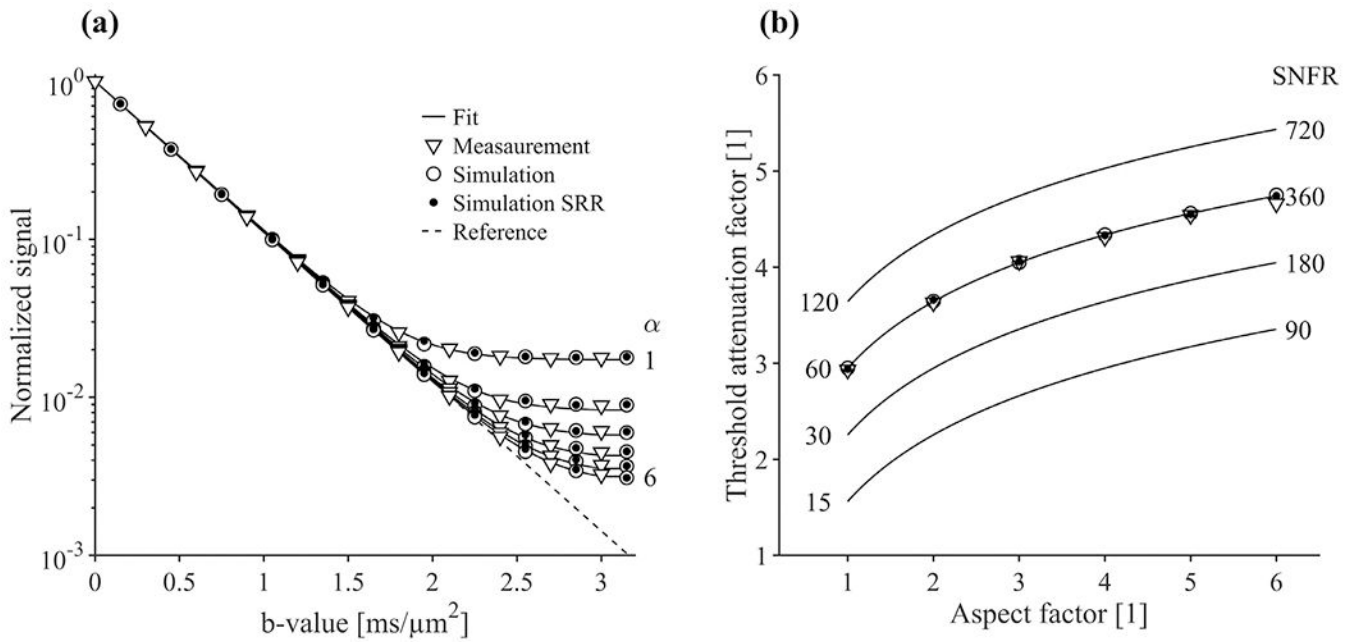
Fig. 2.

Effect of the rectified noise floor on signal accuracy for aspect factors (α) up to 6. Panel (a) shows the simulated and measured mean signal in a water phantom as a function of the b-value, normalized to signal at b0. Panel (b) shows the threshold attenuation factor where the signal bias caused by the noise floor exceeds 5% versus aspect factor. This is shown for measured SNFR (SNFR 60 at α = 1, SNFR = 360 at α = 6), as well as for various SNRF ranging from 15 for α = 1 to 720 for α = 6. The SNFR grows linearly with voxel size, and the threshold attenuation factor shows similar trends for varying SNFR. Sampling with larger voxels improves accuracy and allows for the use of higher b-values before the noise floor affects the signal. Simulations and experimental results (circles versus triangles) show high agreement.
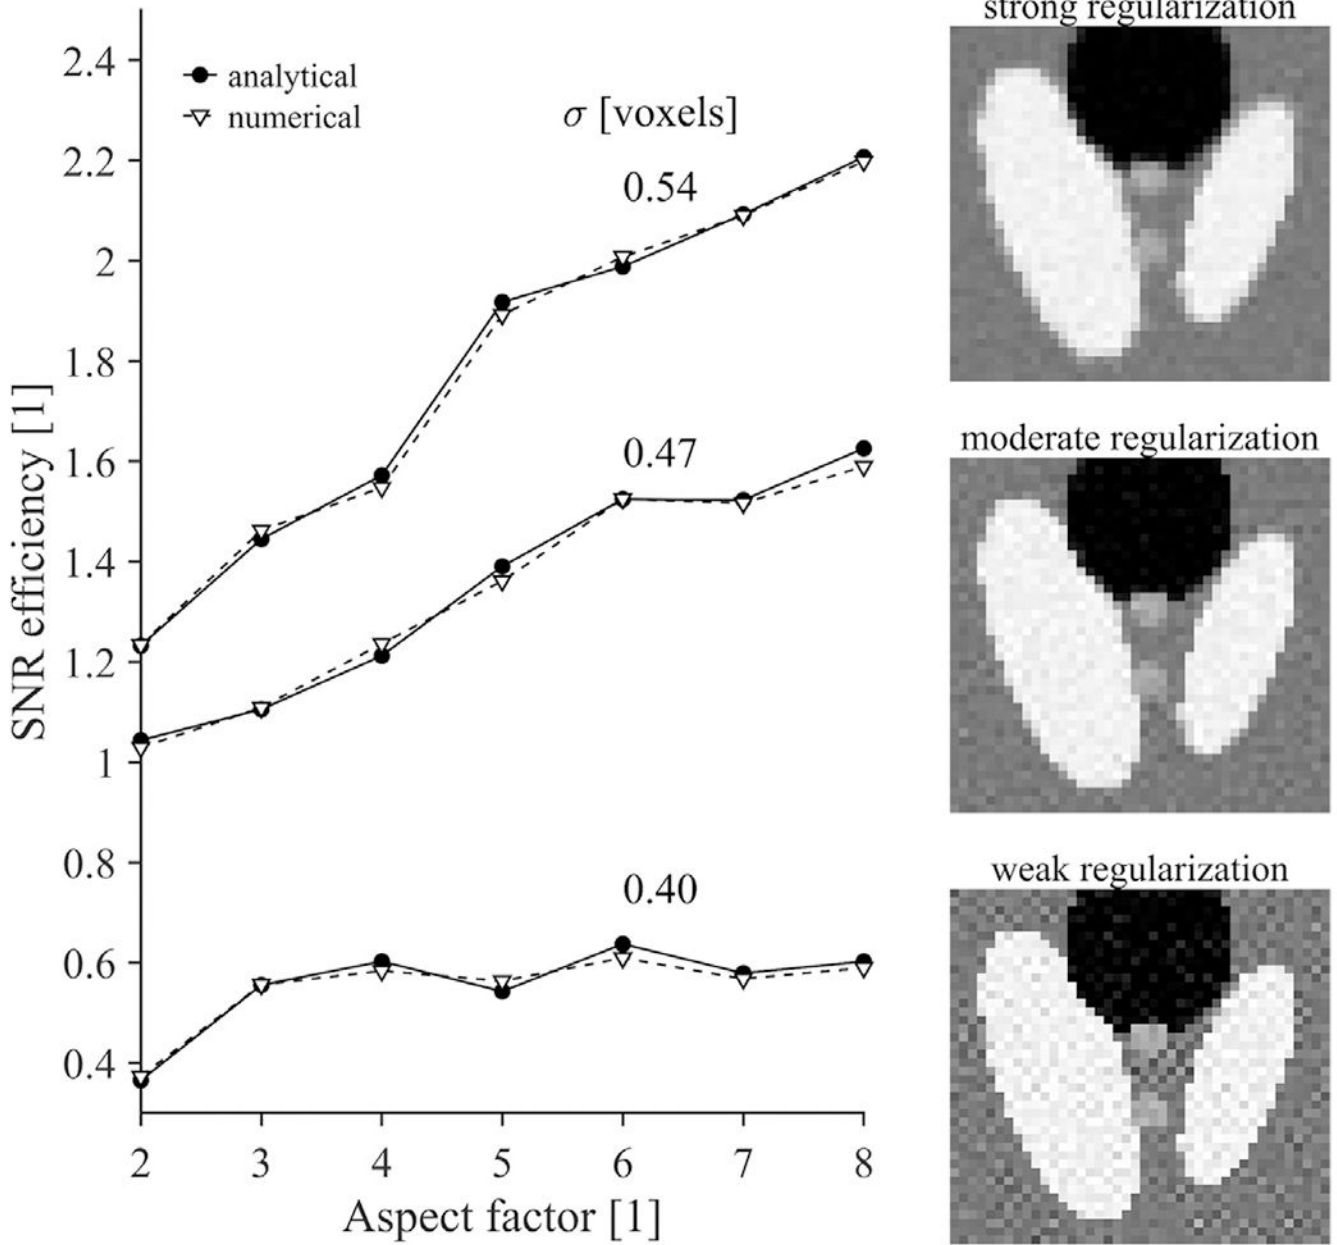
Fig. 3.
SNR efficiency as a function of aspect factor at multiple regularization strengths, evaluated numerically (dashed line) and analytically (solid line). An SNR efficiency above unity means that precision increases compared to direct sampling with matched spatial resolution and acquisition time. Both the increase in aspect factor and regularization strength lead to a higher SNR efficiency. The magnified Shepp-Logan phantom images (right) show the effect of regularization on image characteristics (spatial resolution/noise propagation), and are reconstructed from 5 low-resolution images with an aspect factor of 3 and a SNR of 30 in the region mimicking white matter. Effects of T1-relaxation were not simulated in this analysis.
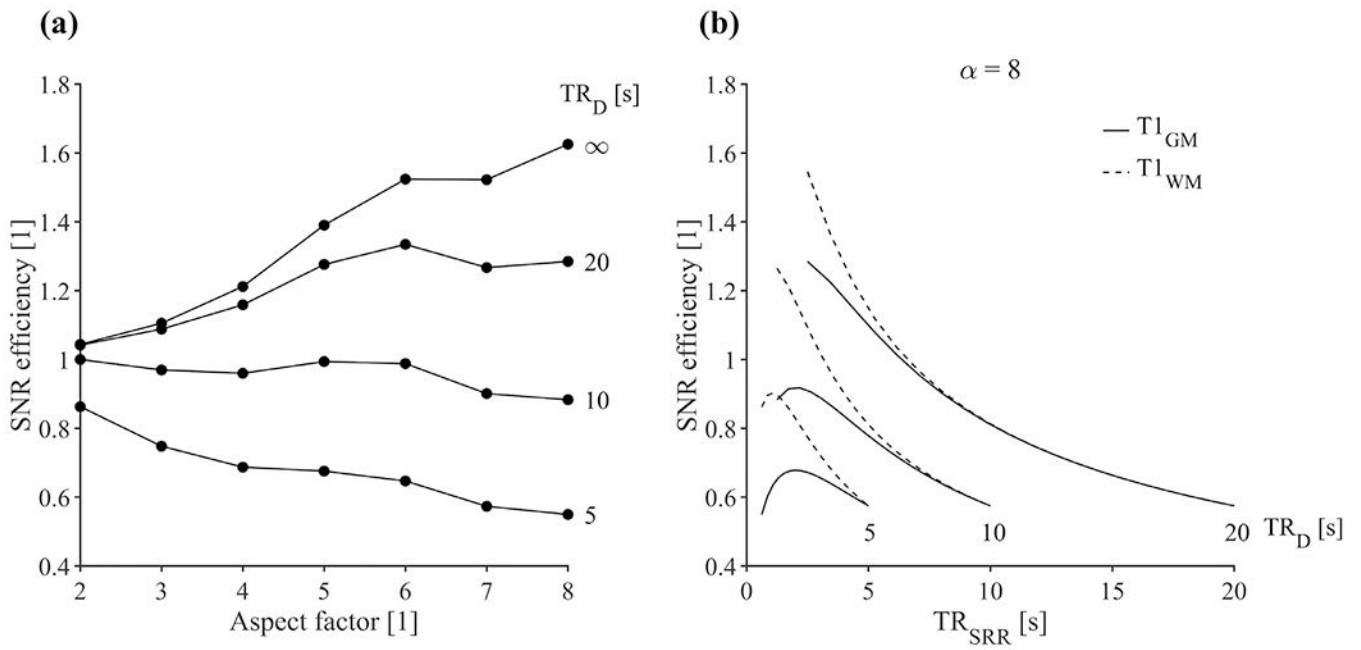
Fig. 4.
Effect of T1 relaxation on SNR efficiency. Panel (a) shows the SNR efficiency using T1 = 1.6 s and the minimum TR_{SRR} for a given TR_D and aspect factor. SNR efficiency decreases as TR_D decreases, with a faster decrease for higher α. Panel (b) shows the effect of changing TR_{SRR} above its minimum at the expense of scan time, for $\alpha = 8$ and T1 of 1.6 s and 0.8 s, mimicking gray- and white matter, respectively. SNR efficiency is maximized when $TR_{SRR}/T1$ is as close as possible to 1.25. Simulations were performed with moderate regularization.
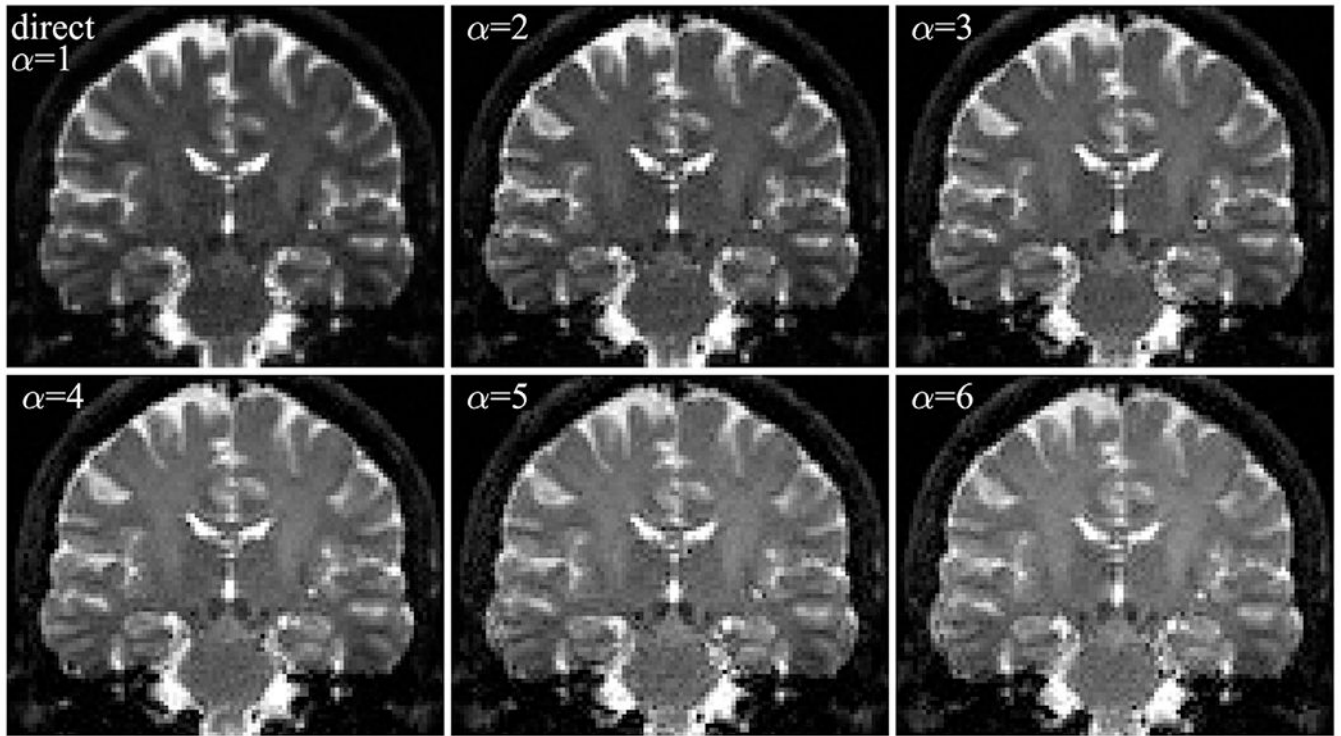
Fig. 5.
Direct sampling and super-resolution reconstruction (SRR) for various aspect factors (α). Image were obtained in a healthy brain in vivo at $b = 0$ ms/μm^2 and reconstructed at a resolution of $1.6 \times 1.6 \times 1.6$ mm^3. All images were reconstructed with moderate regularization and are thus matched in terms of effective resolution ($\sigma = 0.47$ voxels), meaning Gaussian smoothing was applied to the direct case. Due to different signal intervals, the maximum of the colormap is scaled separately for each image. Note that the overall image contrast changes with aspect factor due to T1 effects because a higher aspect factor allows a shorter TR. This results in a smaller signal difference between the white matter (T1 = 0.8 s) and the cerebrospinal fluid (T1 = 4.3 s).
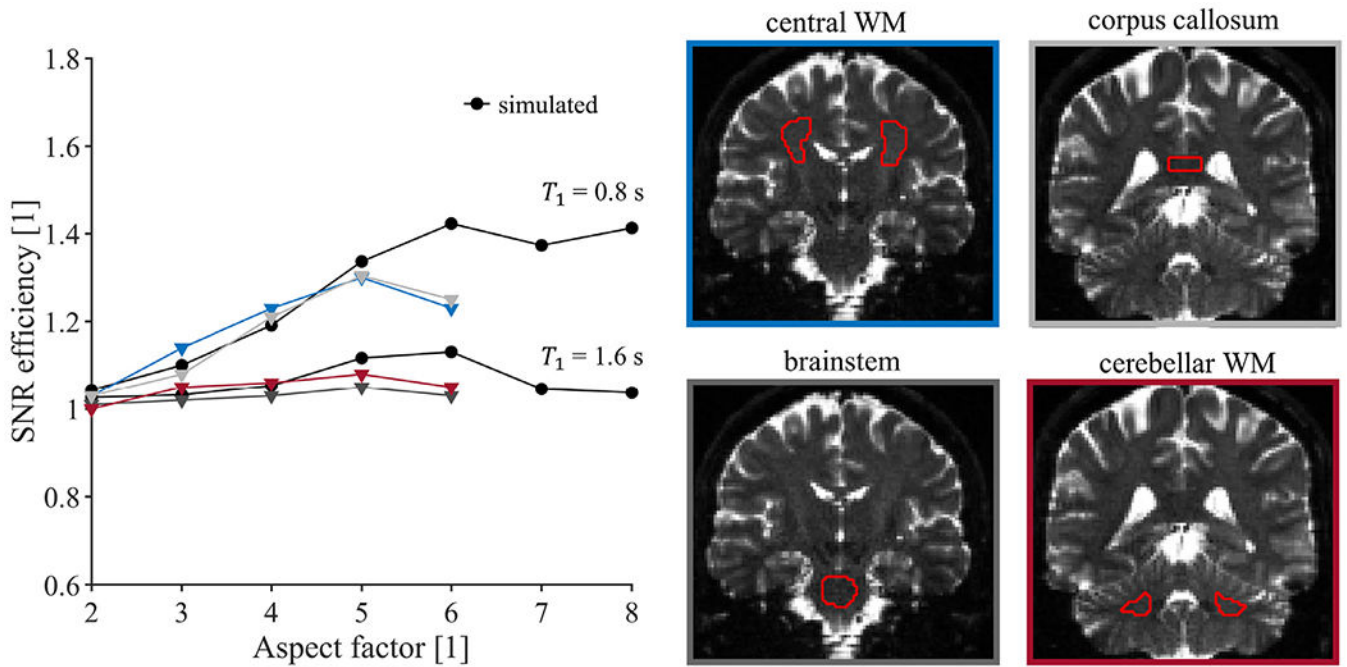
Fig. 6.

SNR efficiency in vivo. The values were obtained from simulations as well as estimated from experimental data, for different aspect factors. In central white matter and the corpus callosum, trends up to $\alpha = 5$ are similar to simulations with T1 = 0.8 s, mimicking white matter. The SNR efficiency in brainstem and cerebellum is more similar to simulations with T1 = 1.6 s, mimicking gray matter. Generally, the SNR efficiency in white matter is overestimated by the simulations (black circles) compared with measurements (triangles). Note that SNR efficiency is above unity in all cases.
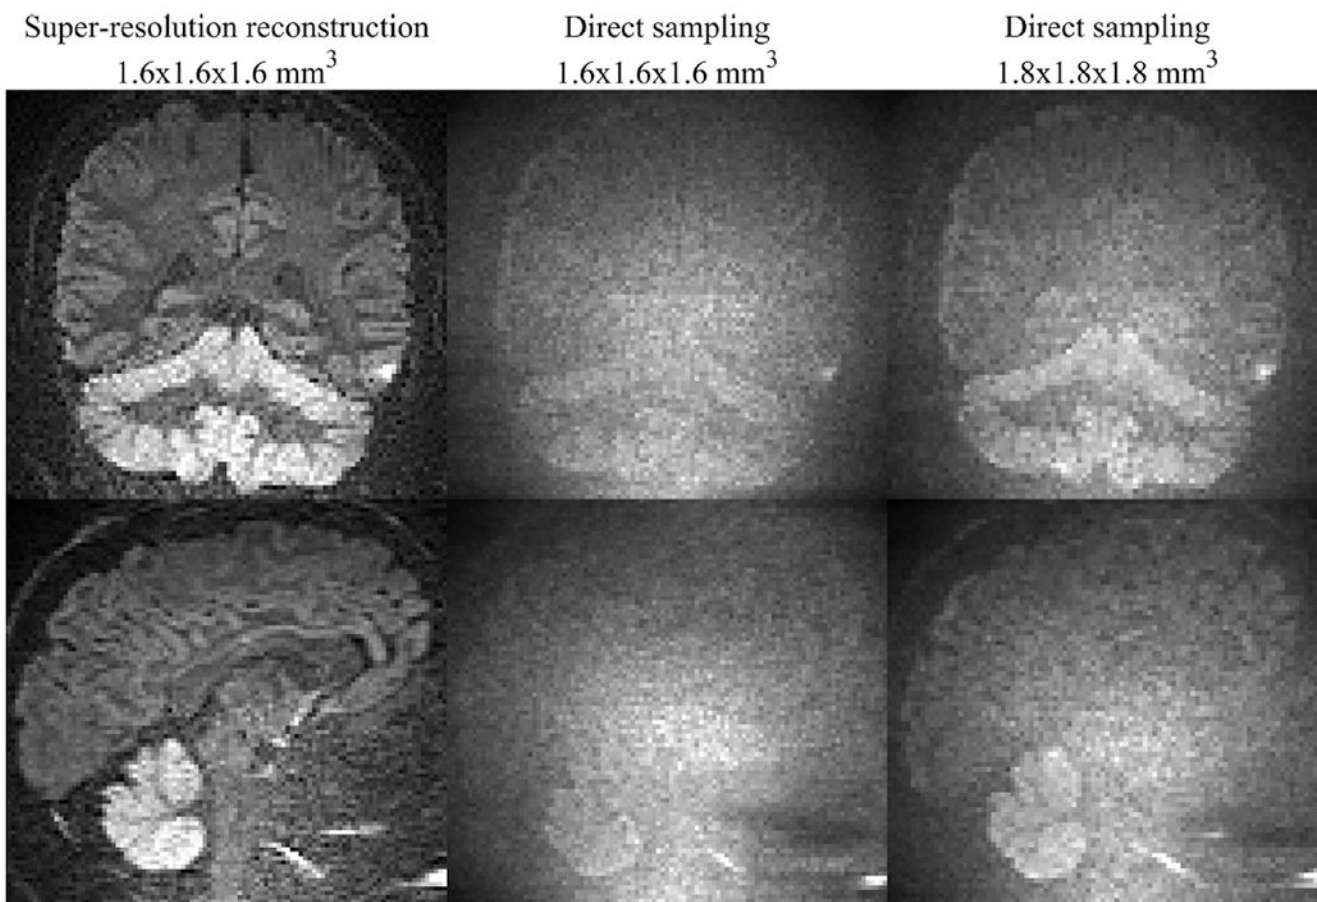
Fig. 7.

In vivo illustration of the accuracy-related benefit of super-resolution reconstruction (SRR). The figure shows diffusion-weighted images from a SRR protocol (left) and direct acquisition (middle) with spherical encoding at $b = 4$ ms/μm^2 in coronal and sagittal view at $1.6 \times 1.6 \times 1.6$ mm^3. As regularization in SRR slightly decreases spatial resolution, we also compare to a direct acquisition at $1.8 \times 1.8 \times 1.8$ mm^3 (right). A vastly higher contrast is observed with SRR compared with direct sampling at both the target resolution and the slightly reduced resolution. Quantitatively, this corresponds to an increase in the contrast ratio between the cerebellar cortex and white matter to 1.86 from 1.06 (middle) or 1.31 (right). Acquisition times are similar.
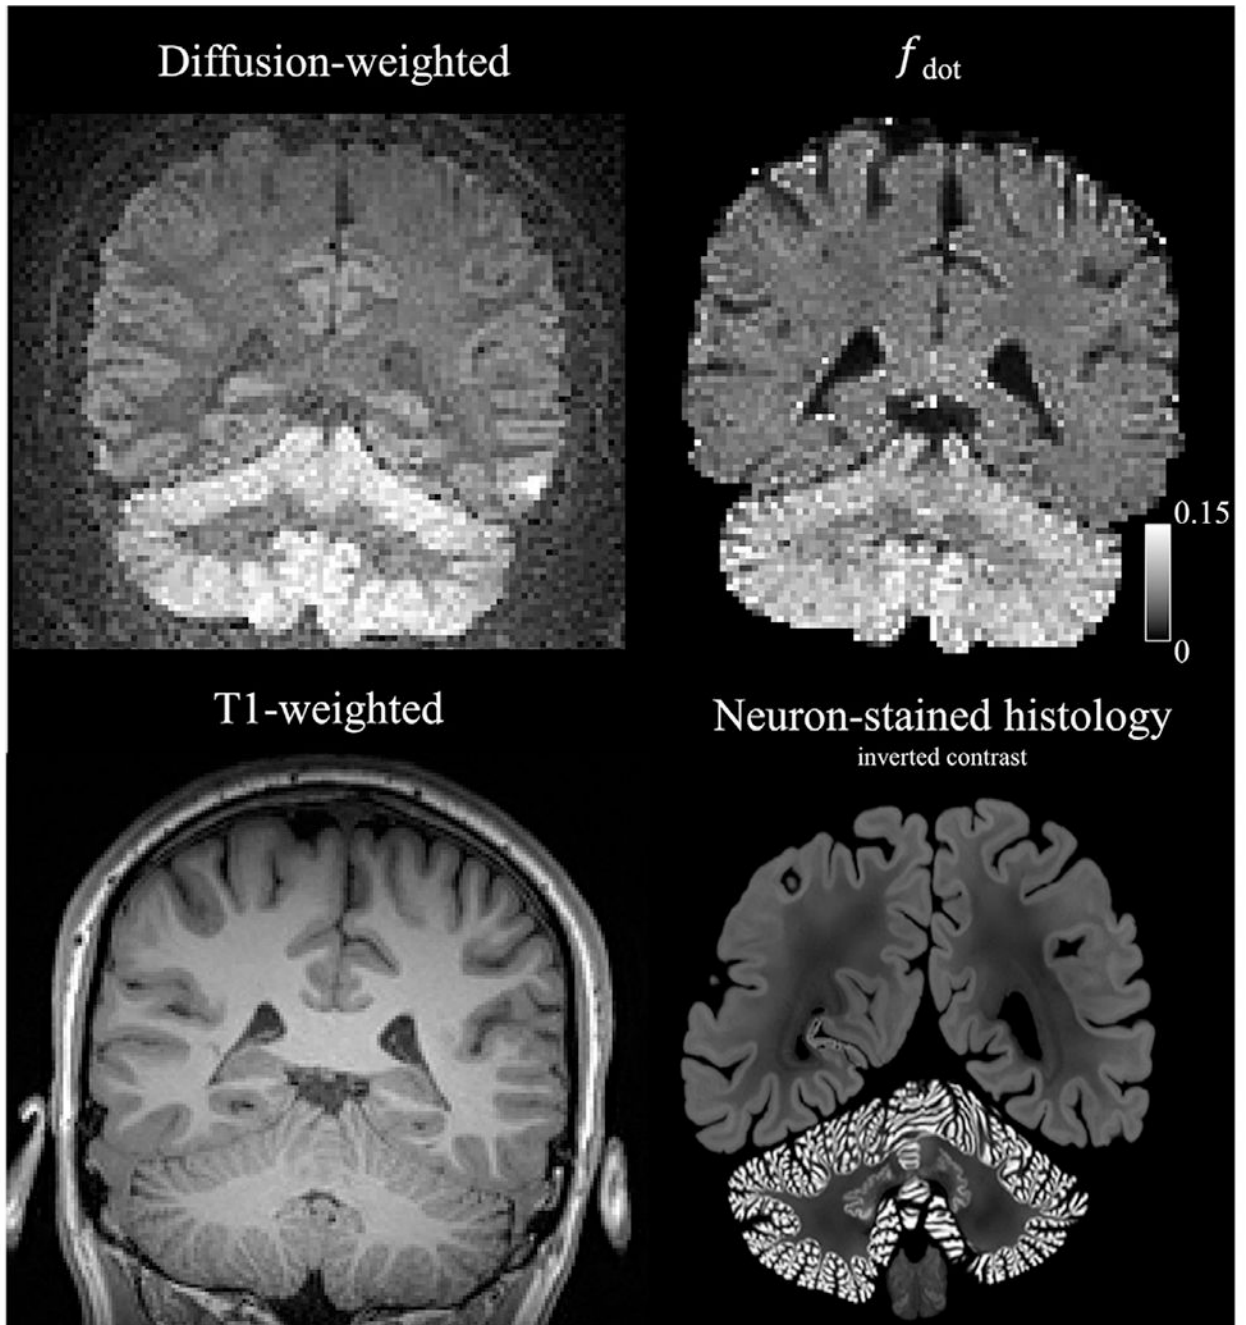
Fig. 8.
Analysis of the image contrast. The top row shows signal retention using diffusion-weighted imaging with spherical encoding at $b = 4$ ms/μm^2 (left) and estimation of f_{dot} (right). The bottom left panel shows a T1-weighted image of the corresponding slice. The bottom right image shows a neuron-stained histology slice from the human brain, obtained from the BigBrain atlas (Amunts et al., 2013). Regions of high diffusion-weighted signal and high f_{dot} are found in the cerebellar cortex where granule cells are densely packed.

Table 1

in vivo acquisition and reconstruction parameters for super-resolution reconstruction protocols. Note that the three numbers that describe the spatial resolution (Res) state the in-plane resolution first, whereas the final value is the slice thickness.

	Direct, α=1	α=2	α=3	α=4	α=5	α=6
N_R [1]	1	3	5	7	8	10
Res [mm³]	$1.6 \times 1.6 \times 1.6$	$1.6 \times 1.6 \times 3.2$	$1.6 \times 1.6 \times 4.8$	$1.6 \times 1.6 \times 6.4$	$1.6 \times 1.6 \times 8.0$	$1.6 \times 1.6 \times 9.6$
TR [ms]	13,000	6500	4300	3300	2600	2200
TA [min]	2:23	3:35	4:00	4:14	3:49	4:02
λ' [1]	–	0.100	0.043	0.028	0.021	0.018