Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Practical design of diffusion MRI based on the decay of spatial modulations in a steady-state

John Gore

john.gore@vanderbilt.edu

Vanderbilt University Medical Center

Bibek Dhakal

Vanderbilt University

Research Article

Keywords: diffusion, longitudinal magnetization, Cramér–Rao bound, protocol optimization, low-field MRI, magnetic resonance microscopy

Posted Date: November 9th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11642008/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Practical design of diffusion MRI based on the decay of spatial modulations in a steady-state

Bibek Dhakal[1,2] and John C. Gore[1–4]
[1] Vanderbilt University Institute of Imaging Science, Vanderbilt University, Nashville, TN, USA [2] Department of Physics and Astronomy, Vanderbilt University, Nashville, TN, USA [3] Department of Biomedical Engineering, Vanderbilt University, Nashville, TN, USA [4] Department of Radiology and Radiological Sciences, Vanderbilt University, Nashville, TN, USA
Correspondence. John C. Gore, Vanderbilt University Institute of Imaging Science, 1161 21st Avenue South, Nashville, TN 37232, USA.
Running title. Optimal design of longitudinal-modulation diffusion MRI
Word count. ~5000 (body, excluding abstract, appendix, references, tables and captions).

Abstract

Purpose: To determine the precision achievable and identify optimal protocols for imaging diffusion using the decay of a spatial modulation of longitudinal magnetization in a spoiled gradient-echo train, for MR microscopy, clinical 3T imaging and low-field systems.
Theory and Methods: A modulation of longitudinal magnetization at spatial frequency *k* decays at a rate $\Lambda = R_1 - \ln(\cos\alpha)/TR + Dk^2$ during a steady state sequence, and D can be obtained from measuring the rates at two spatial frequencies. We derive the constraints that gradient limitations place on image acquisitions and extend the analysis to encoding modules that write several harmonics simultaneously. We calculate the Cramér–Rao bound on D for realistic experimental measurements on typical systems. Protocols were optimized considering modulation period, flip angle, TR, train length and recovery delay at fixed scan time, and checked with Monte Carlo and Bloch–Torrey simulations.
Results: Optimal modulation periods were 65–75 μm at 3 T, near 75 μm at 64 mT and 20–25 μm for microscopy, with shallow optima. Predicted precision was 3.8% per voxel in six minutes at 3 T with a single-shot readout , 13.6% at 64 mT and 18.7% in microscopy, or 0.4–1.9% over a region of 100 voxels. A single 1-1 module is optimal when a recall gradient is needed; a 1-3-3-1 module is better when its harmonics can be resolved. Away from boundaries D was mapped to within 0.3%.
Conclusion: Precise diffusion measurements are feasible with weak gradients on clinical and low-field scanners. The controlling design constraint is that the recall gradient must fit inside TR.

Keywords: diffusion; longitudinal magnetization; Cramér–Rao bound; protocol optimization; low-field MRI; magnetic resonance microscopy

INTRODUCTION

Diffusion-weighted MRI (dMRI) conventionally encodes displacement with pulsed field gradients,[1] and its capabilities are therefore tied to gradient design and performance. The lobes needed to reach a given b value lengthen as the gradient is weakened, and the echo time and the signal lost by T_2 relaxation grow with them. This is a particular constraint for low-field and low-cost systems and for samples with short T_2, while in MR microscopy the gradients of the imaging sequence themselves impose diffusion limits on resolution.[2,3]

We have described an alternative method of dMRI in which diffusion is measured from the decay of a spatial modulation of the longitudinal magnetization M_z during a spoiled gradient-echo train.[4] The modulation is written once with a weak gradient, the diffusion weighting accumulates while the magnetization is stored longitudinally with no gradient applied, and the decay is sampled from a single preparation. Measuring the decay rate at two spatial frequencies gives D independently of relaxation and of the transmit field. The decay law was validated with numerical solutions of the Bloch–Torrey equation, and the method was demonstrated at 15.2 T in phantoms and in plant tissue.[4] It is related to stimulated-echo diffusion measurements,[5] which also store a modulation longitudinally, and to the steady-state transfer function that describes how diffusion of M_z limits resolution in spoiled gradient-echo imaging.[6]

Here we consider how well this method works in different regimes and how measurements should be designed. For any application and system there are several parameters that must be considered including the modulation period, flip angle, TR, number of excitations, recovery delay, readout strategy and form of the preparation module. The choice of each depends on the specific hardware used while the finest usable modulation depends on the gradient system. Here we analyse the precision of measurements with the Cramér–Rao bound and derive optimized protocols for three settings: MR microscopy of fixed tissue at very high field (e.g. 11.7 T), in vivo brain imaging at 3T, and brain imaging at 64 mT on a low-cost, point-of-care system. We also consider four regimes of implementation, from a gradient-free limit in which the structure of the object itself provides the modulation, through modulations resolved in the image, which permit their direct visualization and the simultaneous use of several harmonics in microscopy, to unresolved modulations recovered by a recall gradient in vivo.

THEORY

Decay law and the two-frequency estimate

We first summarize our earlier analysis.[4] Following Torrey,[7] each spatial-frequency component of M_z evolves independently. In a spoiled gradient-echo train with ideal spoiling and uniform D, a modulation at a spatial frequency *k* at which the object itself has no structure decays so that its amplitude before the *n*th pulse is $A_n = A_0 \exp(-\Lambda n TR)$, with

$$\Lambda = R_1 - \frac{\ln(\cos\alpha)}{TR} + Dk^2 = R_{\mathrm{tot}} + Dk^2 \qquad (1)$$

where $R_{tot} = R_1 - \ln(\cos\alpha)/TR$ is the total rate of loss of longitudinal magnetization in the train, comprising relaxation and radiofrequency saturation. R_{tot} is common to all spatial frequencies, so rates measured at two spatial frequencies give

$$D = \frac{\Lambda(k_2) - \Lambda(k_1)}{k_2^2 - k_1^2} \qquad (2)$$

independently of R_1, M_0 and the flip angle, and hence of the transmit field. A convenient choice is to make k_1 small enough that Dk_1^2 is negligible, so that the first measurement serves as a reference for R_{tot} alone.

Writing and recalling the modulation

The modulation can be written with a binomial pulse pair separated by a gradient, the 1-1 module of tagging by spatial modulation of magnetization.[8] Two 90° pulses about the same axis, separated by a gradient of area $\gamma^{-1}k$, give $M_z = -M_0 \cos kx$, a single harmonic with no uniform component (Figure 1b); a DANTE train with a constant gradient is an equivalent alternative.[9] Only the gradient area matters, $\gamma^{-1}k = 23\,500/P$ mT m^{-1} ms for a period P in micrometres (Table 1), and a weak gradient can deliver it by being on for longer though T_2 decay occurs while the magnetization is transverse between the pulses. When the modulation is finer than a voxel it cannot be seen in the image directly, but a gradient of the same area may be applied after each excitation pulse to recall the modulated component to the centre of k space before read-out (Figure 1a). The write gradient is applied once per preparation and its duration is limited in practice by T_2. The recall is applied in every repetition and has to fit inside TR. This single inequality,

$$\gamma^{-1}k \;\le\; G_{\max}(TR - t_{\mathrm{read}}) - \frac{G_{\max}^{\,2}}{S} \qquad (3)$$

in which the last term is the area lost in the ramps of a trapezoidal lobe of slew rate S, sets the finest usable modulation on a given system, and it is what makes practical protocols differ between systems. When the modulation is resolved in the image, as it can be in microscopy, no recall is needed at all and the amplitude is obtained from the Fourier transform of the image itself.

Radiofrequency timing

The flip-angle ratios of the preparation are those of a binomial pulse, but the interval between the pulses is set by the gradient, and it has to have area $\gamma^{-1}k$. For a trapezoidal lobe of amplitude G_{max} and slew rate S, the centre-to-centre spacing of the two 90° pulses of a 1-1 module is

$$\tau_{12} = t_p + \frac{\gamma^{-1}k}{G_{\max}} + \frac{G_{\max}}{S} \qquad (4)$$

where t_p is the duration of a hard pulse, the plateau being shortened to $\gamma^{-1}k/G_{max} - G_{max}/S$, because the two ramps contribute G_{max}^2/S of area. Table 1 lists the spacing for each of the systems considered here. At 3T, with 40 mT m^{-1} and a slew rate of 150 T m^{-1} s^{-1}, a 65 µm modulation needs 9.8 ms between the two pulses; at 10 mT m^{-1} and 20 T m^{-1} s^{-1} the same modulation needs 37.6 ms, and in both cases the gradient rather than the radiofrequency accounts for almost all of it. A higher-order module repeats the interval at the base spatial frequency once for each gap between RF pulses, so a 1-3-3-1 module lasts about three times the spacing listed for its base period.

After the preparation module, a 1-1 pair of flip angle α produces a pattern $M_z = M_0(\cos^2\alpha - \sin^2\alpha \cdot \cos kx)$. An error in the transmit-field reduces the modulation only as $\sin^2\alpha$ and the remainder becomes a uniform term at k = 0, where it does not interfere with the measurement. For example, a 30% error in B_1 reduces the amplitude by 21% but

introduces no bias in estimates of D as the amplitude is a free parameter of the fit as described below. In addition, a resonance offset Δf accumulates a phase $2\pi\Delta f\cdot\tau_{12}$ between the pulses, and that phase adds to kx. The modulation is translated by $2\pi\Delta f\cdot\tau_{12}/k$ but their amplitude is unchanged. Long pulse spacings are thus tolerable even where the field is poorly shimmed, and the same is true of higher-order modules, in which every harmonic is displaced by the same shift of the pattern.

Phase alternation

Mapping D requires the uniform component of M_z — the steady state of the train, the term relaxing toward it, and the unmodulated fraction $\cos^2\alpha_p$ left by imperfect preparation pulses of flip angle α_p — to be separated from the modulation. Repeating the preparation with the phase of its final pulse advanced by 180°, as in CSPAMM,[10] writes +cos *kx* instead of −cos *kx*, and the complex difference of the two image series contains only the modulation, which decays at Λ voxel by voxel independently of R_1, M_0 and the accuracy of the preparation pulses.[4] The dual preparation costs nothing in signal-to-noise ratio per unit time compared with averaging two images but the subtraction must be made on complex images, and the two preparations interleaved so that bulk motion is common to the pair.

Several spatial frequencies from one preparation

Higher-order binomial pulses sharpen the stripes of the modulations by adding harmonics, and those harmonics may be used to advantage. For example, a 1-3-3-1 module writes amplitudes of 0.930, 0.389 and 0.070 at k, 2k and 3k (Figure 1b). Each harmonic is an independent spatial-frequency mode of the Bloch–Torrey equation and decays at its own rate, so a single preparation yields

$$A_m(n) = a_m\, e^{-\left(R_{\mathrm{tot}}+Dm^2k^2\right)nTR}, \qquad m = 1,2,3 \qquad (5)$$

with the diffusion weighting of the *m*th harmonic m^2 times that of the fundamental.
Equation (5) suggests the data may be analyzed in more than one way to estimate D. First, the amplitudes may be fitted harmonic by harmonic, by log-linear regression of ln A_m against *n*, and the resulting rates Λ_m regressed against m^2k^2. The slope is D and the intercept R_{tot}, which is Equation (2) generalized. Alternatively a single weighted nonlinear least-squares may be fit to Equation (5) for all harmonics simultaneously, with R_{tot} and D shared and the amplitudes free. Third, where a map rather than a regional value is wanted, the model may be fit voxel by voxel with D constrained to vary smoothly, so that the precision of neighbouring voxels is pooled without blurring the boundaries of the map. In each case a coarse modulation, for which Dk^2 is negligible, may be included as a reference. This fixes R_{tot} and improves the estimation of D. The choice of module depends on whether a recall gradient is needed as summarized in Table 2.

The effective diffusion weighting

The amplitude of the modulation is attenuated by $\exp(-Dk^2t)$ during an interval t in which it is stored longitudinally. The magnetization is also reduced during the gradient lobes during which it is transverse, by $k^2\tau_w/3$ for each lobe (see below), so the effective b value of image n is

$$b_{\mathrm{eff}}(n) = k^2(nTR + \ell\tau_w/3) \qquad (6)$$

where ℓ is the number of lobes (see Equation (7)). The lobe term is the same in every image, so it scales the amplitude but does not affect the fitted rate or D. For the realistic systems considered below it is a small contribution to the total at 3T (6%) and in microscopy (7%), but higher (23%) at 64 mT, where the gradients are weak and the lobes are longer. The diffusion weighting of importance accumulates without gradients. The gradient required to write the modulation has area $\gamma^{-1}k$; for a period of 100 μm this is 235 mT m^{-1} ms, for example 24 mT m^{-1} for 10 ms. This modulation survives for a time of order $1/(R_{tot} + Dk^2)$, which sets the useful range of t.

No gradients at all

If the object itself contains structure at a spatial frequency where the modulation is measurable, no preparation gradient is needed. Our recent paper gives the steady-state transfer function for a spoiled sequence, from which D may in principle be recovered by comparing images acquired with different values of R_{tot}[6] such as using different flip angles. The information is integral to the object's own edges, but is confined to within a few λ (λ = $\sqrt{(D/R_{tot})}$ is the distance water diffuses during the mean age of 1/Rtot of the magnetization) of them, so the method is only practical when the voxel dimension is comparable with λ. This condition may be met in magnetic resonance microscopy but not in vivo, as discussed below.

Signal model

The signal-to-noise-normalized amplitude of image *n* is

$$A(n) = c\, S_1 \sin\alpha\ f_{rec}\, e^{-\ell\tau_w/T_2}\, e^{-\ell D k^2 \tau_w/3}\, e^{-\Lambda n TR} \qquad (7)$$

where S_1 is the signal-to-noise ratio of one fully relaxed image at the stated voxel size and readout bandwidth, τ_w is the duration of the write gradient, $f_{rec} = 1 - (1 - M_{z,ss}/M_0)\exp(-T_{rec}/T_1)$ is the longitudinal magnetization recovered in the delay T_{rec} between the end of one train and the next preparation, assuming the train reaches the spoiled steady state $M_{z,ss}$.

The factors in τ_w are transverse relaxation and diffusion while the magnetization is in the transverse plane under a gradient lobe. During a lobe the wavenumber of the phase grows linearly between 0 and k, so each lobe contributes $b = k^2\tau_w/3$. ℓ is the number of such lobes.

When the modulation is resolved in the image, it can be read without a recall gradient, and c = 1, ℓ = 1. When it is not resolved, every image needs a recall lobe equal to the write lobe, so ℓ = 2. The recall also brings to the centre of k-space only one of the two components $e^{\pm ikx}$ that make up cos kx, which gives c = ½, as for a stimulated echo. The two phase-alternated images carry ±A(n) with independent noise and their half-difference has amplitude A(n) and noise reduced by √2.

METHODS

Precision and protocol optimization

To estimate achievable precision and design practical protocols (Tables 2–4), we considered fixed-tissue microscopy at 11.7 T, in vivo brain imaging at 3 T, and brain imaging at 64 mT on a low-cost, point-of-care system. Table 3 gives the assumed relaxation properties, diffusivities, and gradient limits. We calculated precision from the

Cramér–Rao bound for a joint fit to the decay curves, treating each initial amplitude as a free parameter while sharing Rtot and D across curves. Monte Carlo simulations of the same estimator agreed with the bound to within 1.6%. The amplitude of each image was given by Equation (7).
Protocols were optimized considering the modulation period, flip angle, TR, number of excitations per preparation and recovery delay. The constraints were Equation (3) and a fixed total scan time, which sets the number of acquisitions or phase-encoding steps available. Precision is quoted per voxel so for a region of V voxels it improves as $\sqrt{V}$.

Mapping when D varies with position

To test whether D can be mapped when it varies with position, the longitudinal magnetization was propagated pulse by pulse by solving the Bloch–Torrey equation in one dimension on a periodic grid of 0.5 μm, applying the propagator for diffusion and relaxation exactly over each repetition period and multiplying by cos α at each excitation,[4] with D varying in steps between slabs (Figure 3). Voxel-wise D was obtained from the difference between the decay rates at the modulation spatial frequency and at k = 0, using only images whose modulation exceeded the noise of the protocol.

Simulated brain slice

To compare the bound with the scatter of a complete reconstruction, and to examine how the modulation period affects voxels containing cerebrospinal fluid (CSF), we used the simulated axial brain slice described previously:[4] anatomy from the IXI dataset,[11] 2 mm voxels, white matter, cortical and deep grey matter and CSF with their own relaxation times, proton density and D, the 3 T protocol of Table 3 with a six-minute acquisition, and complex noise at a single-image signal-to-noise ratio of 54. D was fitted in every voxel from the noise-corrected power against the k = 0 reference.

RESULTS

Choice of modulation period

Table 4 and Figure 2 give the predicted precision as a function of the modulation period for each of the three protocols. The precision varies little near its best value, which lies at 65–75 μm at 3 T, near 75 μm at 64 mT and at 20–25 μm for microscopy, where the much shorter T_1 and the larger R_{tot} of a higher flip angle require a correspondingly larger Dk^2. The optimum is set by a balance of three effects: a finer modulation increases Dk^2 relative to R_{tot}, but it shortens the modulation lifetime so that fewer excitations contribute, and it demands a longer recall gradient, which lengthens TR. Where the recall cannot be fitted into TR at the available gradient amplitude the entry is omitted; this is what excludes periods below about 50 μm on the low-field system.

Protocols for microscopy, 3 T and low field

Table 3 sets out a complete protocol for each setting, chosen under the constraint of Equation (3). The three are remarkably similar in the quantity that matters. All use a modulation period between 20 and 70 μm, and the two in vivo protocols require almost the same gradient area, 361 and 340 mT m^{-1} ms. However, the time taken to deliver it —

9 ms at 40 mT m^{-1} against 34 ms at 10 mT m^{-1} — and thus the shortest TR achievable differ between gradient systems of different capabilities.
Readout strategy strongly influences precision. If one phase-encoding line is acquired per excitation, each image of the series needs its own preparation for every line, and the available scan time is spent on phase encoding. If instead a single-shot readout is used, the whole series of images follows one preparation and the scan time is spent on averaging. At 3 T, for a six-minute acquisition, a segmented acquisition of 64 lines gives a predicted 13% per voxel while a single-shot readout, with its per-image signal-to-noise ratio reduced by a factor of about three for the wider bandwidth, gives 3.8%. The single-shot form is also free of the sensitivity to motion between shots that afflicts any tagging measurement. The useful train is short — ten excitations at 3 T, five at 64 mT, twenty in microscopy — because the modulation lifetime is 10–100 ms, and most of each shot is the recovery delay needed to restore M_z before the next preparation.

Low-field and low-cost systems

The requirement that the preparation imposes on the gradient system is one of area, not amplitude, and the diffusion weighting then accumulates with no gradient applied. With a maximum gradient of 10 mT m^{-1}, a 69 μm modulation needs 34 ms of gradient to write and the same again to recall before every read-out. TR becomes 64 ms, which is in any case a natural choice for a low-field readout, and the predicted precision is 13.6% per voxel in six minutes at a single-image signal-to-noise ratio of 20, or 1.4% over a region of a hundred voxels (Table 3). An effective b value of 640 s mm^{-2} accumulates in the 78 ms lifetime of the modulation, during which no gradient is applied. A further 190 s mm^{-2} is carried by the write and recall lobes themselves (Equation 6). Per image, the 68 ms of transverse evolution in the write and recall lobes and the factor ½ of the recall leave 19% of the amplitude. A pulsed-gradient experiment reaching the same b value with the same 10 mT m^{-1} would need two 51 ms lobes and an echo time of at least 103 ms, and would retain 28% for T_2 = 80 ms. Per image the two are comparable at 10 mT m^{-1}, where the ratio is largest (0.77), and the tagged signal becomes the smaller of the two at weaker gradients (0.51 at 3 mT m^{-1}, 0.30 at 2 mT m^{-1}). The benefit of the preparation at low field is therefore not signal per image. It is that the diffusion weighting is set by gradient area and accumulates without gradient, and that several images are read from each preparation.

Choice of preparation module

Table 2 compares binomial preparation modules. When the modulation is unresolved, each harmonic needs its own recall gradient, of area *m* times the write area, so at a fixed peak gradient area a higher-order module places its highest harmonic at the frequency a 1-1 module would use and divides the excitations among coarser harmonics. A single 1-1 module at the highest affordable spatial frequency is then the most precise, 3.8% at 3 T against 7.6% and 13.2% for two- and three-harmonic modules (Table 2a). When the modulation is resolved, no recall is needed and all harmonics are read simultaneously from the same images. In the microscopy setting the 1-3-3-1 module is then better than the 1-1 module at equal write-step area, by a factor of 1.4–1.9 that grows as the available step area decreases (Table 2b).

Mapping D

Figure 3 shows D recovered voxel by voxel from a sample in which D varies in steps. With the in vivo protocol of Table 3, a 65 μm modulation and 2 mm voxels, the five voxels of uniform D return 0.698, 0.701, 0.999, 1.400 and 1.400 against true values of 0.70, 1.00 and 1.40 μm² ms⁻¹, errors below 0.3%. The voxel straddling a step returns 1.150 against its volume-weighted value of 1.20. With the microimaging protocol, a 20 μm modulation and 15 μm voxels, voxels more than two modulation periods (45 μm) from a step are accurate to 0.21%, and those between one and two periods also to 0.21%. Voxels at a step are less accurate, by an amount that depends on where the step falls within the voxel and on the direction of the change. At the rise from 0.45 to 0.75 μm² ms⁻¹ the voxel containing the step reads 4.1% low and the next voxel 0.7% high; at the fall from 0.75 to 0.60 μm² ms⁻¹ the corresponding errors are 0.5% and 1.1% low. All are within the 19% per-voxel precision of that protocol.

Precision in a simulated brain slice

In the simulated brain slice,[4] the voxel scatter of D in voxels that contained at least 90% one tissue was 5.1–5.6% for white matter, deep grey matter and cortical grey matter in six minutes. With each tissue's own relaxation times, proton density and D, the Cramér–Rao bound for this protocol is 4.1–4.6%. The difference arises from residual partial volume in voxels that are 90% rather than 100% one tissue.

The modulation period must also be matched to the range of D. With a 65 μm modulation the pattern in CSF decays within one TR (Dk^2 = 28 s^{-1}, lifetime 33 ms), so 219 of 225 fluid voxels could not be measured, whereas a 120 μm modulation returned 2.98 ± 0.23 μm² ms⁻¹ for fluid, with a median error of 6.3% over the brain. In voxels containing 30–50% fluid the volume-weighted D is 1.86 μm² ms⁻¹, but the 120 μm map returned 1.50, an under-estimate of 19%, and the 65 μm map 1.28. Two effects weight the fit towards tissue. The long T_1 of fluid recovers only about a third of M_0 in the 1.25 s delay, so that per unit volume fluid contributes about 60% of the signal of tissue, and its modulation decays about twice as fast (lifetime ≈ 90 ms against ≈ 180 ms), so that its share of the signal falls along the train and the single-exponential fit follows the slower tissue component in the later images. The partial-volume bias therefore grows as the period is shortened, and the period that is optimal for the precision of tissue values is not optimal where fluid fractions matter.

DISCUSSION

The precision of this measurement is governed by gradient area rather than gradient amplitude, and by the number of images read from each preparation. The optimal modulation period balances the diffusion weighting, which favours fine modulations, against the lifetime of the modulation and the length of the recall gradient, which favour coarse ones. The optima are shallow, so that a period near 65–75 μm serves both 3 T and 64 mT in vivo, and the much shorter T_1 and larger R_{tot} of fixed tissue at high field move the microscopy optimum to 20–25 μm. Readout strategy matters more than any single timing parameter: a single-shot readout, by reading the whole decay from one preparation, improves precision by a factor of about three over a segmented acquisition at 3 T, and is also free of the sensitivity to motion between shots.

The four regimes. With no gradients at all, information is derived from the object's own structure and is confined to within a few λ of its boundaries. λ cannot exceed $\sqrt{(D \cdot T_1)}$, which is ≈19 μm for fixed tissue and ≈34 μm for brain in vivo, so this regime requires voxels no larger than a few tens of micrometres and is restricted to microscopy. With a weak preparation gradient and a resolved modulation, which requires voxels smaller than about half the modulation period, D is mapped at a resolution set by the modulation period rather than by the voxel, a natural regime for microscopy and for high-resolution animal imaging. With an unresolved modulation the amplitude must be recovered by a rewinding gradient before readout, which removes the resolution requirement and is the appropriate implementation in vivo.

The resolved regime has two further advantages. The result is a map of D at the voxel resolution while the displacement probed is set by the modulation period, so the two are decoupled, and as no recall gradient is used, higher-order modules deliver several spatial frequencies from one preparation, which improves precision substantially (Table 2b).

Low-field imaging. The scaling that governs conventional diffusion encoding is difficult to reach on low cost, low field systems. To reach a given b value the gradient lobe must lengthen and it increases the echo time with it, so the signal lost to T_2 grows rapidly as the gradient is weakened. The preparation described here is governed by area alone, and is followed by a gradient-free interval in which the diffusion weighting continues to accumulate, so that it need not affect the echo time. Its duration τ_w lengthens as G^{-1}, and when the modulation is unresolved it is incurred twice, in the write and in the recall, while the recall refocuses only half of the modulation. Per image the tagged signal is therefore within a factor of about 0.8 of a pulsed-gradient echo of equal b at 10 mT m^{-1} and falls below it at weaker gradients. Its advantages lie in the gradient-free diffusion interval and in reading many images from each preparation. On a 64 mT system limited to 10 mT m^{-1} the predicted precision is only 3.6 times poorer than at 3 T with 40 mT m^{-1}. A factor of 2.7 comes from the lower image signal-to-noise ratio and the remainder mainly from the longer write and recall lobes. Low field is also favourable in two incidental respects: the static field inhomogeneity expressed in hertz is small, so T_2^* approaches T_2 and the long gradient costs little extra; and the long readouts that low field imposes leave ample room inside TR for the recall gradient. Against this, the short T_1 of tissue at low field limits the accessible diffusion time to a few hundred milliseconds.

Matching the reference. Experiments at 15.2 T[4] identified a mismatch between the reference and modulated preparations as the main source of systematic error when Dk^2 is small relative to R_{tot}. Effects that act on the modulated series but not on the reference, such as eddy currents of the gradient lobes or incompletely spoiled stimulated-echo pathways, appear directly as an error in D. Therefore the reference should match the modulated preparation in every respect except the spatial frequency, and Dk^2 should be at least comparable to R_{tot}. The protocols of Table 3 use a coarse reference with the same timing and lobes and have $Dk^2/R_{tot} = 2.4$ in microscopy and at 3 T; the 64 mT protocol, at 1.1, sits at the edge of this condition and is the one most affected by a mismatched reference.

Limitations. Only the first regime is entirely free of gradients, and it is the least precise. The preparation gradients needed in the others are weak and brief but still essential. Bulk motion during the measurement may displace the modulation. In the resolved regime this appears as a shift that can be measured and corrected, but in the unresolved regime it

causes signal loss, as it does in stimulated-echo experiments. The unresolved regime also costs signal even without motion. The recall gradient refocuses only one of the two components $e^{\pm ikx}$ that make up the stored cos kx, so each image carries half of the modulation (c = ½ in Equation 7), as in a stimulated echo, and the magnetization spends both the write and the recall lobe in the transverse plane. Together these roughly double the predicted error in D compared with a resolved read-out of the same modulation. Part of the lost half could be recovered by reading the conjugate echo after a further recall lobe of twice the area. That echo, however, carries much more diffusion weighting of the first during the recall and a further $2\tau_w$ of T_2 decay, so at 3 T it returns only about 0.55 of the first echo, and combining the two improves the signal-to-noise ratio by a factor of about 1.14 rather than $\sqrt{2}$, at the cost of a longer TR. The loss is avoided altogether where the voxel is small enough to resolve the modulation. Coherent flow washes the modulation out and will be read as an elevated apparent diffusivity, which may be a confound or, at long diffusion times, a measurement of perfusion.[12] In tissue with more than one compartment the decay of the modulation is weighted toward long-T_1 components, and exchange between compartments during the long observation times will influence the measured value; this is a property of the measurement rather than an artefact. Finally, the analysis assumes ideal spoiling.[13] With imperfect spoiling, stimulated-echo pathways return transverse coherence to the longitudinal axis and the simple geometric decay of the modulation acquires further terms, which can be suppressed by adequate spoiling or by phase cycling.

The predicted precision of Tables 2–4 follows from the Cramér–Rao bound and assumes that the model is correct. Monte Carlo simulations of the estimator agreed with the bound to within 1.6%, and the reconstruction of a simulated brain slice came within about 25% of it. The decay law, the reference subtraction and the tensor inversion have been shown to behave as predicted on a real system,[4] but the in vivo protocols proposed here remain to be implemented.

CONCLUSIONS

Diffusion measurements based on the decay of a longitudinal modulation in a steady-state train can be designed for systems ranging from microscopy to ultra-low field. The finest usable modulation is set by the requirement that the recall gradient fit inside TR; the optimal periods are 20–25 μm in microscopy and 65–75 μm in vivo; and a single 1-1 module read by single-shot imaging is preferred when the modulation is unresolved, whereas multi-harmonic modules are advantageous when it is resolved. Optimized protocols are predicted to give about 4% per voxel at 3 T and 14–19% per voxel at 64 mT and in microscopy, or 0.4–1.9% over a region of a hundred voxels, at realistic signal-to-noise ratios.

APPENDIX: ACQUISITION AND ANALYSIS FOR IN VIVO MAPPING

The steps of a suggested measurement at 3 T are as follows. Four image series are acquired, being the modulated and reference preparations each acquired with two tag phases.

1. Wait a recovery delay of 1.25 s, which restores 68% of M_0.

2. Apply the tagging along the chosen axis: a hard 90°$_x$ pulse, a gradient lobe of area 361 mT m^{-1} ms (40 mT m^{-1} for 9 ms), and a second hard 90°$_x$ pulse 9.8 ms after the first. Follow it with a spoiler.
3. Apply the imaging train: ten excitations of 25°, TR = 40 ms, each comprising a recall gradient of 361 mT m^{-1} ms along the modulation axis followed by a single-shot read-out, so that image *n* is acquired at t = nTR.
4. Repeat steps 1 to 3 with the phase of the second 90° pulse advanced by 180°, which writes +cos kx.
5. Repeat steps 1 to 4 with the write and recall gradients switched off, which leaves M_z inverted uniformly and gives the reference series at k = 0, the limiting case of the coarse modulation assumed in Table 3.
6. Repeat the whole of steps 1 to 5 for averaging; fifty-four repetitions occupy six minutes.

The analysis proceeds voxel by voxel.

1. Subtract the two phase-alternated series as complex images, pairing acquisitions that are adjacent in time so that any bulk displacement is common to both. Take the magnitude of the difference only after the subtraction, and average over repetitions in magnitude rather than in the complex plane: a rigid displacement δ advances the phase of the recalled modulation by kδ, which is a full radian for δ = 10 μm, but leaves its magnitude unchanged.
2. The recall gradient places the modulated component at the centre of k space, so the magnitude of each voxel of image *n* is the local amplitude A(n) and no demodulation is needed. Where the modulation is resolved instead, band-pass the image about +k and take the magnitude of the complex result.
3. Average the *squared* magnitudes rather than the magnitudes themselves. The mean square is $\langle |z|^2 \rangle = A^2 + 2\sigma^2$, with σ the noise of one component of the difference image, so subtracting the measured noise variance leaves A^2 exactly, at every signal-to-noise ratio; averaging the magnitudes themselves leaves a residual floor that flattens the decay and biases D high by a few percent. Fit $\ln(\langle |z|^2 \rangle - 2\sigma^2)$ against *n* weighted by the corrected power, or fit A(n) = a exp(−ΛnTR) by weighted nonlinear least squares with a free amplitude.
4. Form D = [Λ(k) − Λ(0)]/k², with k = 2π/65 μm, in each voxel. Equivalently fit the two series jointly with R_{tot} and D shared and the two amplitudes free.
5. Repeat steps 1 to 4 with the write and recall gradients set to zero, keeping the same RF pulses and timing. This leaves M_z uniformly inverted and gives the reference series at k = 0, the limiting case of the coarse modulation assumed in Table 3. A coarse modulation written and recalled with lobes of the same timing is an alternative that also shares the gradient lobes, provided its period is not resolved in the image; if it is resolved, the reference amplitude has nodes at which its rate cannot be measured, and these appear as stripes in the map of D. The result is the diffusion coefficient along the modulation axis. A tensor requires the preparation to be repeated along at least six directions, the gradient area being distributed between the axes so that the total area, and hence the pulse spacing, is unchanged.

REFERENCES

1. Stejskal EO, Tanner JE. Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient. *J Chem Phys.* 1965;42:288–292.
2. Kaiser R, Bartholdi E, Ernst RR. Diffusion and field-gradient effects in NMR Fourier spectroscopy. *J Chem Phys.* 1974;60:2966–2979.
3. Callaghan PT. *Principles of Nuclear Magnetic Resonance Microscopy.* Oxford: Clarendon Press; 1991.
4. Dhakal B, Gore JC. Measuring diffusion from the decay of longitudinal magnetization patterns in steady-state MRI. *Sci Adv.* Submitted.
5. Tanner JE. Use of the stimulated echo in NMR diffusion studies. *J Chem Phys.* 1970;52:2523–2526.
6. Dhakal B, Gore JC. Diffusion broadening of the point spread function in steady-state MRI. *Magn Reson Med.* Submitted.
7. Torrey HC. Bloch equations with diffusion terms. *Phys Rev.* 1956;104:563–565.
8. Axel L, Dougherty L. MR imaging of motion with spatial modulation of magnetization. *Radiology.* 1989;171:841–845.
9. Mosher TJ, Smith MB. A DANTE tagging sequence for the evaluation of translational sample motion. *Magn Reson Med.* 1990;15:334–339.
10. Fischer SE, McKinnon GC, Maier SE, Boesiger P. Improved myocardial tagging contrast. *Magn Reson Med.* 1993;30:191–200.
11. IXI Dataset. Information eXtraction from Images. Imperial College London. https://brain-development.org/ixi-dataset/ (accessed 23 October 2026). Licensed under CC BY-SA 3.0.
12. Le Bihan D, Breton E, Lallemand D, Grenier P, Cabanis E, Laval-Jeantet M. MR imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders. *Radiology.* 1986;161:401–407.
13. Zur Y, Wood ML, Neuringer LJ. Spoiling of transverse magnetization in steady-state sequences. *Magn Reson Med.* 1991;21:251–263.

TABLES

TABLE 1 Gradient area and radiofrequency pulse spacing needed to write a modulation of a given period. The area $\gamma^{-1}k$ is a property of the modulation alone; the spacing is the centre-to-centre interval between the two 90° pulses of a 1-1 module, Equation (4), evaluated at the maximum gradient, slew rate and hard-pulse duration of each system: 1000 mT m^{-1}, 5000 T m^{-1} s^{-1} and 0.1 ms for microscopy; 40 mT m^{-1}, 150 T m^{-1} s^{-1} and 0.5 ms at 3 T; 10 mT m^{-1}, 20 T m^{-1} s^{-1} and 1.0 ms at 64 mT. A 1-3-3-1 module repeats the interval at its base period three times, so it lasts about three times the value listed. When the modulation is unresolved the same area must also be delivered as a recall gradient inside every TR, which is the constraint of Equation (3).

Modulation period (μm)	Gradient area $\gamma^{-1}k$ (mT m^{-1} ms)	Pulse spacing, MR microscopy (ms)	Pulse spacing, 3 T (ms)	Pulse spacing, 64 mT (ms)
200	117	0.42	3.7	13.2
150	157	0.46	4.7	17.2
100	235	0.53	6.6	25.0
75	313	0.61	8.6	32.8
65	361	0.66	9.8	37.6
50	470	0.77	12.5	48.5
35	671	0.97	17.5	68.6
25	939	1.24	24.3	95.5
20	1174	1.47	30.1	118.9

TABLE 2 Choice of preparation module. (a) When the modulation is unresolved every harmonic needs its own recall gradient, of area *m* times the write area, so modules are compared at equal peak area; the 3 T protocol of Table 3 is used throughout. A single 1-1 module at the highest affordable spatial frequency is the most precise, because the higher-order modules place little amplitude at that frequency and divide the excitations between harmonics. (b) When the modulation is resolved no recall is used, the harmonics are read simultaneously, and modules are compared at equal write-step area. Resolving periods of 37–224 μm requires voxels smaller than half the shortest period, so the microscopy protocol of Table 3 (15 μm voxels) is used, with flip angle, TR (≥ 10 ms) and train length re-optimized for every entry. Here the three-harmonic module is the better choice, by a margin that grows as the available step area decreases. Precision is quoted per voxel. Panels (a) and (b) use different protocols (3 T and microscopy) and should not be compared with each other; each compares modules within one protocol.

Module	Base period (μm)	Spatial frequencies measured (μm)	σ(D)/D
(a) Modulation unresolved; peak gradient area 361 mT m^{-1} ms			
1-1	65	65	3.8%
1-2-1	130	130, 65	7.6%
1-3-3-1, two harmonics	130	130, 65	7.6%
1-3-3-1, three harmonics	195	195, 98, 65	13.2%

Module	Base period (μm)	Spatial frequencies measured (μm)	σ(D)/D
(b) Modulation resolved, no recall; equal write-step area			
1-1, step area 105 mT m^{-1} ms	224	224	113%
1-3-3-1, step area 105 mT m^{-1} ms	224	224, 112, 75	58%
1-1, step area 157 mT m^{-1} ms	150	150	52%
1-3-3-1, step area 157 mT m^{-1} ms	150	150, 75, 50	31%
1-1, step area 209 mT m^{-1} ms	112	112	31%
1-3-3-1, step area 209 mT m^{-1} ms	112	112, 56, 37	22%

TABLE 3 Protocols for three settings, chosen under the constraint that the recall gradient fits inside TR at the available gradient amplitude (Equation 3). All three use a coarse reference modulation, of period 2 mm for microscopy and 4 mm in vivo, in a second preparation of equal duration; the scan time covers both. None of the modulations is resolved in the image, so the amplitude follows Equation (7) with c = ½ and ℓ = 2: the magnetization spends both the write and the recall lobe in the transverse plane, and the recall refocuses half of the modulation. Precision is the Cramér–Rao bound for the fit described in Methods, quoted per voxel and for a region of a hundred voxels, at the stated signal-to-noise ratio of a single fully relaxed image. b_{eff} includes the weighting of the write and recall lobes, Equation (6).

	MR microscopy	In vivo brain	In vivo brain
Field	11.7 T, fixed tissue	3 T	64 mT
T_1 / T_2 (ms)	500 / 30	1430 / 70	250 / 80
D assumed (μm² ms^{-1})	0.7	0.8	0.8
Maximum gradient (mT m^{-1})	1000	40	10
Voxel	15 μm	2.0 mm	2.5 mm
Readout	one line per excitation	single shot	single shot
Modulation period (μm)	20	65	69
Gradient area, write and recall (mT m^{-1} ms)	1174	361	340
Duration at maximum gradient (ms)	1.2	9.0	34
TR (ms)	10	40	64
Flip angle (°)	40	25	30
Excitations per preparation	20	10	5
Recovery delay (ms)	1900	1250	200
R_{tot} (s^{-1})	28.7	3.2	6.2
Dk^2 (s^{-1})	69.1	7.5	6.6
Dk^2/R_{tot}	2.4	2.4	1.1
Modulation lifetime 1/Λ (ms)	10	94	78
b_{eff} at t = 1/Λ (s mm^{-2})	1087	935	832
Amplitude surviving the write and recall module	0.88	0.74	0.37

	MR microscopy	In vivo brain	In vivo brain
Longitudinal recovery factor f_{rec}	0.98	0.68	0.86
Repetitions in the scan	128 phase-encoding steps	109 averaged	346 averaged
Scan time, both preparations (s)	540	360	360
Single-image SNR assumed	60	54	20
σ(D)/D, single voxel	**18.7%**	**3.8%**	**13.6%**
σ(D)/D, 100-voxel region	**1.9%**	**0.38%**	**1.4%**

TABLE 4 Predicted precision in D against the period of the fine modulation, for each protocol of Table 3 with everything else held at the values given there. An entry is omitted where the recall gradient of that period cannot be fitted into TR at the available gradient amplitude. The microscopy protocol uses a 40° flip angle and therefore has a large R_{tot}, which is why coarse modulations are of little use to it; its precision is best at 20–25 μm. The in vivo protocols have shallow optima. At 3 T the optimum is at 65–75 μm, and the precision is within 13% of its best between 50 and 75 μm and within 20% up to 100 μm. At 64 mT the optimum is near 75 μm, and the precision is within 9% of its best between 65 and 100 μm. At 50 μm it is 65% worse, because the write and recall lobes, 2 × 47 ms at 10 mT m^{-1}, cost much of the amplitude to T_2.

Modulation period (μm)	Gradient area (mT m^{-1} ms)	MR microscopy	In vivo, 3 T	In vivo, 64 mT
200	117	429%	12.6%	33.6%
150	157	244%	7.8%	21.7%
100	235	113%	4.6%	14.0%
75	313	67%	3.9%	13.1%
65	361	53%	3.8%	14.3%
50	470	35%	4.3%	21.7%
35	671	23%	7.0%	—
25	939	18.8%	22%	—
20	1174	18.7%	133%	—

FIGURES

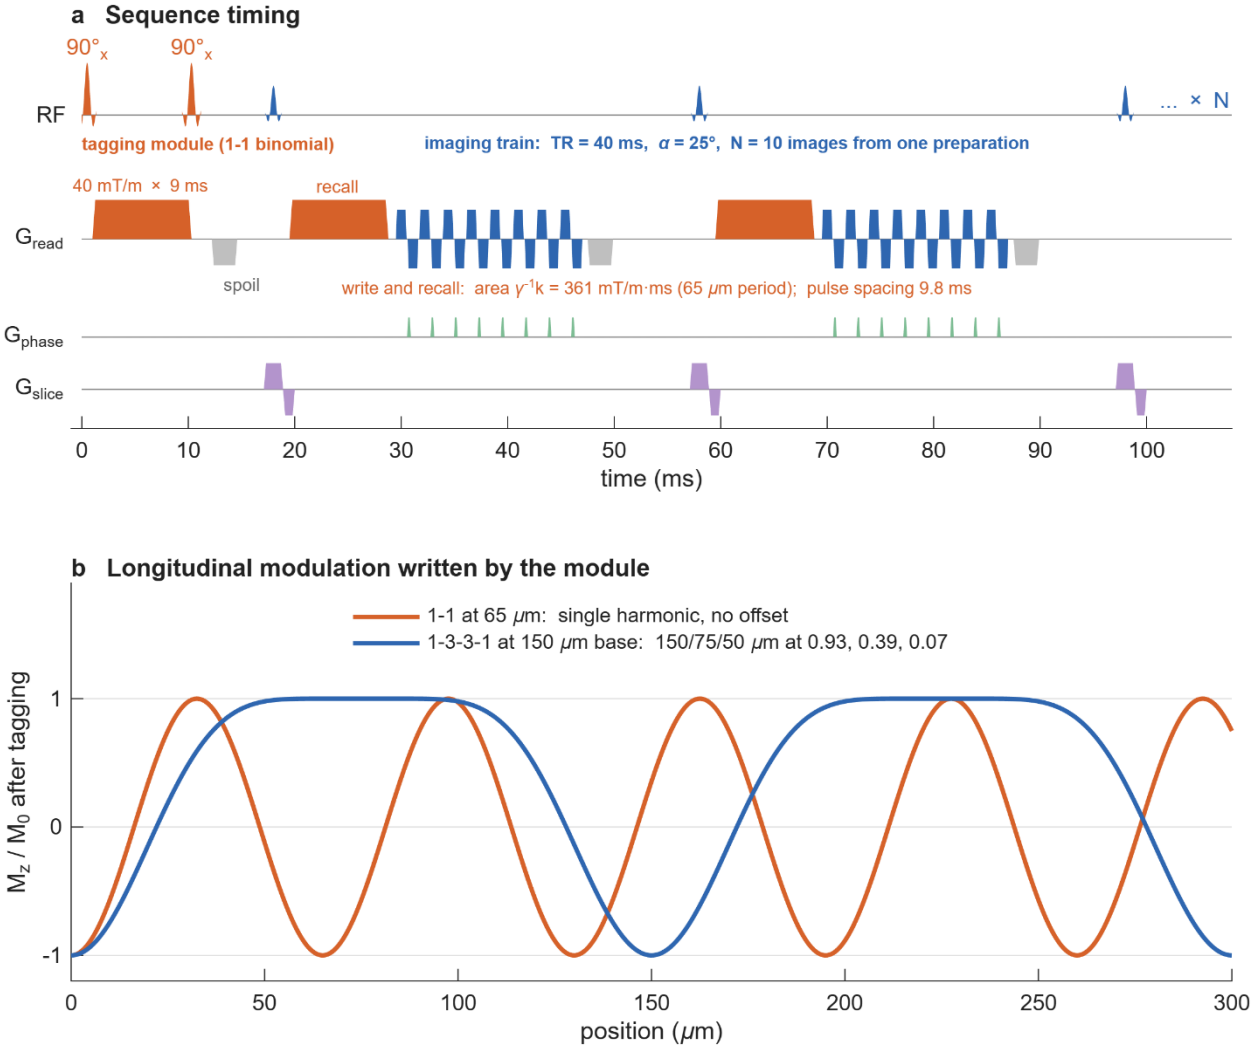


FIGURE 1 Preparation and read-out, drawn for the 3 T protocol of Table 3. (a) A 1-1 binomial module ($90°_x$ – gradient – $90°_x$) writes a sinusoidal modulation of the longitudinal magnetization with a period of 65 µm. The write gradient carries the area $\gamma^{-1}k$ = 361 mT m^{-1} ms, here 40 mT m^{-1} for 9 ms, giving a centre-to-centre pulse spacing of 9.8 ms, and is applied once; a spoiler then removes residual transverse magnetization. The imaging train follows with TR = 40 ms, α = 25° and a single-shot readout, so that one complete image of the series is acquired per excitation. A 65 µm modulation cannot be resolved in vivo, so a recall gradient of the same area precedes every readout to bring the modulated component to the centre of k space; it is this gradient, repeated in every TR, and not the write gradient, that sets the hardware requirement. The second preparation repeats the module with a much smaller area to measure R_{tot} alone. (b) Modulation written by the module. The 1-1 pair with two 90° pulses produces a single harmonic at *k* and no residual uniform component. A 1-3-3-1 module at a 150 µm base produces harmonics at 150, 75 and 50 µm with amplitudes 0.93, 0.39 and 0.07, together with a uniform offset of 0.39, and so measures three spatial frequencies from one preparation; it is useful where the harmonics are resolved and can be read simultaneously (Table 2b).

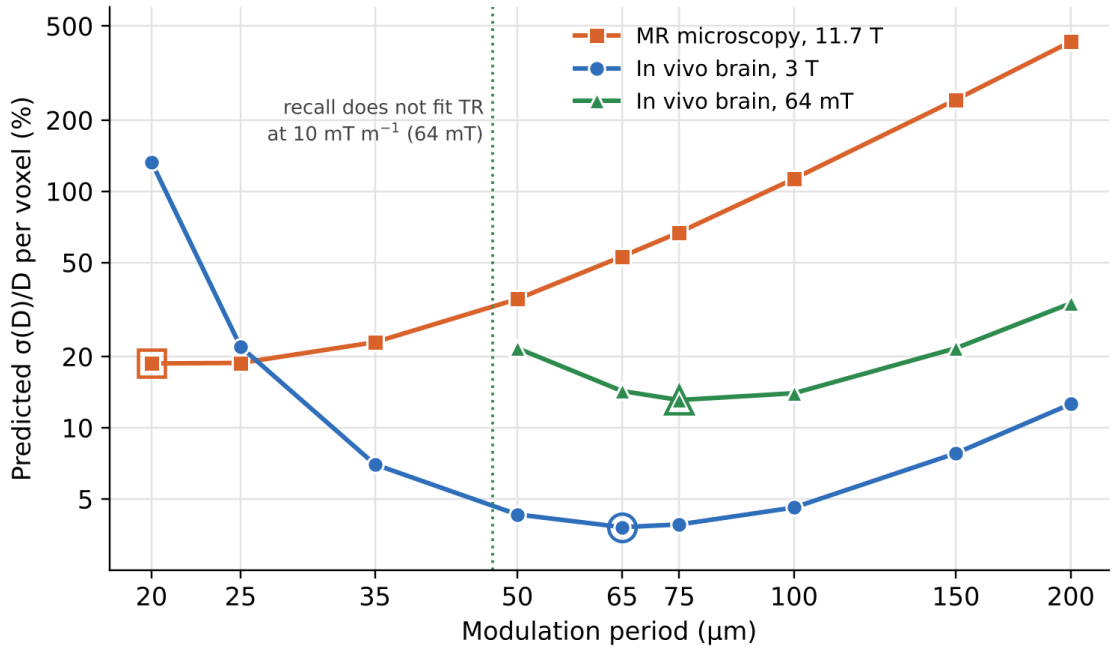


FIGURE 2 Predicted precision of D per voxel against the period of the fine modulation, for the three protocols of Table 3 with everything else held at the values given there (data of Table 4). Open symbols mark the optimum of each protocol. The 64 mT curve ends at 50 μm because finer modulations need a recall gradient that cannot be fitted into TR at 10 mT m^{-1}. Note the logarithmic axes: the in vivo optima are shallow, whereas the microscopy protocol, with its 40° flip angle and large R_{tot}, needs periods of 20–25 μm.

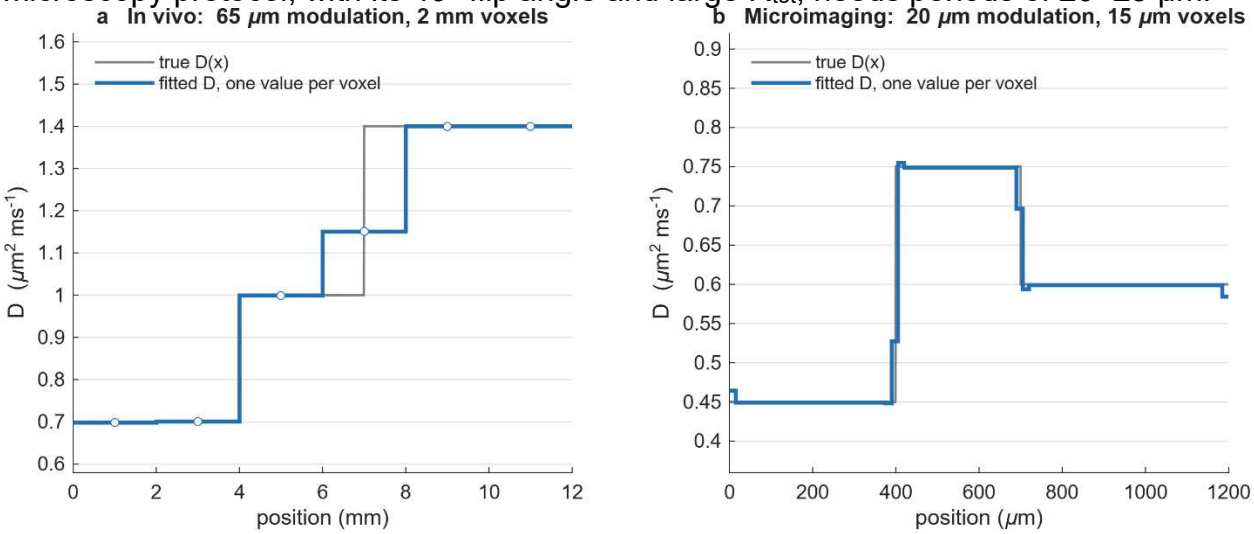


Figure 3 Mapping D when it varies with position, from the pulse-by-pulse solution of the Bloch–Torrey equation with D varying between slabs. The grey line is the true D(x). The blue staircase is the value fitted independently in each voxel from the difference between the decay rates at the modulation spatial frequency and at k = 0. The calculation is noise-free; the noise enters only by restricting the fit to images whose modulation exceeds the noise of the protocol. (a) In vivo protocol of Table 3: 65 μm modulation, 2 mm voxels, TR = 40 ms, α = 25°, ten excitations. Voxels of uniform D are accurate to 0.3%; the voxel straddling a step returns 1.15 against its volume-weighted value of 1.20. (b) Microimaging

protocol: 20 µm modulation, 15 µm voxels, TR = 10 ms, α = 40°, twenty excitations. The map is accurate to 0.2% more than one modulation period from a step, and voxels at a step are in error by at most 4%.