

Sampling of non-Gaussian Ensemble Average Propagators for the simulation of diffusion magnetic resonance images

Justino R. Rodríguez-Galván¹  | Pablo Villacorta-Aylagas¹ |
 Susana Merino-Caviedes¹ | Federico Simmross-Wattenberg¹ |
 Carlos Castillo-Passi²  | Pablo Irarrazaval²  | Antonio Tristán-Vega^{1,3}  |
 Carlos Alberola-López^{1,3} 

¹Laboratorio de Procesado de Imagen (LPI), Universidad de Valladolid, Valladolid, Spain

²Pontificia Universidad Católica de Chile, Santiago, Chile

³Instituto de Investigación Biosanitaria de Valladolid, Valladolid, Spain

Correspondence

Justino R. Rodríguez-Galván, Laboratorio de Procesado de Imagen (LPI), Universidad de Valladolid, ETSI de Telecomunicación, Campus Miguel Delibes, 47011 Valladolid, Spain.
 Email: jroldgal@lpi.tel.uva.es

Funding information

Agencia Estatal de Investigación, Grant/Award Numbers: RTI2018-094569-B-I00, PID2020-115339RB-I00, PID2021-124407NB-I00; ESAOTE Ltd., Grant/Award Number: 18IQBM

Abstract

Purpose: (a) To design a methodology for drawing random samples of any Ensemble Average Propagator (EAP) (b) to modify the KomaMRI simulator to accommodate them as realistic spin movements to simulate diffusion MRI (dMRI) and (c) to compare these simulations with those based on the Diffusion Tensor (DT) model.

Theory and Methods: The rejection method is used for random sampling of EAPs: starting from a probability law that is easily sampled, and whose density function *wraps* the target EAP, samples are accepted when they lie inside the targeted region. This is used to sample the EAP as described by Mean Apparent Propagator MRI (MAP-MRI) and in Spherical Convolution (SC) based on Spherical Harmonics (SH). With this methodology, MAP-MRI and SC representations are calculated over *in-vitro* pig hearts images, and a simulation of a pulsed-gradient spin echo (PGSE) dMRI sequence inside the myocardial wall is undertaken with the KomaMRI simulator.

Results: MAP-MRI shows better agreement with the actual acquisition than conventional DT-based simulations, in terms of Mean Squared Errors and correlation with improvements up to 1.7 % for the former and 2.2 % for the latter.

Conclusion: dMRI sequences can be simulated accurately (yet, efficiently) if phantoms with a proper per-spin description of the diffusion process are made available. Moreover, our findings suggest that the study of non-Gaussian diffusion of the heart might be feasible, at least *in vitro*.

KEYWORDS

diffusion MRI, Ensemble Average Propagator, MRI simulators, random sampling

1 | INTRODUCTION

MRI simulators constitute a useful tool for educational and research purposes. For the former, clinical demands and equipment cost imply extensive scanner use, so hands-on sessions for practitioners can hardly be afforded.^{1,2} For the latter, simulations allow users to isolate specific phenomena by selectively removing undesired effects, such as off-resonance and non-ideal hardware behavior. Consequently, several MRI simulators have been proposed in the literature since the early days^{3,4} and different computationally efficient simulators are available for different purposes,^{5–12} with some including diffusion.^{6,9,13}

Among the diversity of MRI sequences, diffusion MRI (dMRI) has an unparalleled potential to reveal the anisotropy of soft tissues (e.g., the organization of myelinated axons of the white matter of the brain or the muscular fibers of the myocardial wall). This is achieved, among others,^{14,15} by combining the original Spin Echo (SE) sequence with two strong pulsed gradients with the same polarity, plus an inversion pulse in between. This way, the random movement of water molecules during the lapse between the two gradients produces a Diffusion Weighted Image (DWI) whose attenuation depends on the movement direction relative to the orientation of the gradients.¹⁶ Although applications based on Diffusion Tensor Imaging (DTI) implicitly model this movement as Gaussian,¹⁷ DTI is an oversimplified model in many situations.¹⁸

This issue carries over to dMRI simulation: while Gaussian diffusion can be simulated by numerically solving the Bloch-Torrey's equation,¹⁹ more general diffusion is challenging. Alternatively, diffusion can be simulated by tuning the proton density (PD) with a diffusion-driven attenuation model⁹; however, this is only an approximation. Microstructural simulation,^{6,13} on the other hand, relies heavily on an accurate model for the simulated anatomy.

In this work, we aim at simulating non-Gaussian behaviors in dMRI. Instead of attempting to generalize and solve the Bloch-Torrey's equation, we (1) design a solution to draw random samples from general 3D EAPs to simulate spin motion during the diffusion time, and (2) extend the capabilities of the KomaMRI simulator¹¹ to accommodate this motion. To our knowledge, only one procedure for a similar sampling has been used for a diffusion-related task,²⁰ but not for modeling end-to-end spin displacements intended for DWI simulation. We base the description of spin movements on signal representations rather than on biologically inspired models¹⁸ to attain an anatomy-agnostic methodology. Specifically, we have focused on two such representations, namely: Mean Apparent Propagator MRI —MAP-MRI—²¹ and

Spherical Convolution —SC— based on Spherical Harmonics —SH—. ²² However, our approach extrapolates to any other signal representation.

As a proof of concept, we have chosen to simulate diffusion in the myocardium of an *in-vitro* heart.²³ reported that high-order diffusion models, compared to DTI, were unlikely to provide more accurate descriptions due to long TEs together with the T2 signal loss effect. Simulation provides a ground to test the validity of this statement in ideal conditions with a still heart. If in these conditions differences are not observed, the statement seems well founded. Otherwise, room for further research exists when these conditions are not ideal.

2 | THEORY

2.1 | The Ensemble Average Propagator in dMRI

Let δ be the duration of the pulsed gradients and Δ the inter-pulse lapse of the sequence, so that the effective diffusion time is $\tau = \Delta - \delta/3$. Within τ , water molecules are driven by Brownian motion, and their net displacement $\mathbf{R} \in \mathbb{R}^3$ can be modeled as a 3-D random variable (RV) whose Probability Density Function (PDF), $P(\mathbf{R}; \tau, D)$, is dubbed Ensemble Average Propagator (EAP).¹⁶ For Gaussian diffusion:

$$P(\mathbf{R}; \tau, D) = \frac{1}{\sqrt{(4\pi\tau)^3 |D|}} \exp\left(-\frac{\mathbf{R}^T D^{-1} \mathbf{R}}{4\tau}\right), \quad (1)$$

where D is the diffusion tensor (DT), the 3×3 symmetric, positive-semidefinite covariance matrix of the process. Drawing samples from such distribution is straightforward by generating 3 uncorrelated Gaussians with zero mean and unit standard deviation, $\mathbf{S} = [u, v, w]^T : \{u, v, w\} \sim \mathcal{N}(0, 1)$, and computing: $\mathbf{R} = D^{1/2} \mathbf{S}$. Non-Gaussian diffusion does not admit a general expression for the EAP, since $P(\mathbf{R}; \tau, D)$ can be any non-negative, unit-mass distribution that fulfills antipodal symmetry: $P(\mathbf{R}; \tau, D) = P(-\mathbf{R}; \tau, D)$. Among the many methods aimed at describing arbitrary EAP representations, we will focus on the following two:

- Using **Cartesian** coordinates, the MAP-MRI formalism draws the EAP as a basis functions expansion in terms of Hermite's polynomials H_n (with physicists' convention):²¹

$$P(\mathbf{R}'; \tau, D) = \sum_{\substack{N=0 \\ (\text{even})}}^{N_{\max}} \sum_{\substack{n_1+n_2+n_3=N \\ [n_1, n_2, n_3]}} a_{n_1, n_2, n_3} H_{n_1}(x'; s_x) H_{n_2}(y'; s_y) H_{n_3}(z'; s_z) \\ : H_n(t; s) = \frac{1}{\sqrt{2^{n+1}/\pi} n! s} \exp\left(-\frac{t^2}{2s^2}\right) H_n\left(\frac{t}{s}\right), \quad (2)$$

with rotated coordinates $[\mathbf{x}', \mathbf{y}', \mathbf{z}']^T$ and scaling factors $\{s_x, s_y, s_z\}$ related to the spectrum of D as:

$$\mathbf{R}' = U^T \mathbf{R}; \begin{bmatrix} s_x \\ s_y \\ s_z \end{bmatrix} = \begin{bmatrix} \sqrt{2\tau\lambda_x} \\ \sqrt{2\tau\lambda_y} \\ \sqrt{2\tau\lambda_z} \end{bmatrix} : D = U \Lambda U^T, \quad (3)$$

with U the matrix of eigenvectors of D and $\lambda_x \geq \lambda_y \geq \lambda_z$ its eigenvalues, hence $\Lambda = \text{diag}(\lambda_x, \lambda_y, \lambda_z)$.

- Using **Spherical** coordinates, the EAP can otherwise be expressed as the SC of a fiber-Orientation Distribution Function (fODF), $\Phi(\mathbf{r})$, $\mathbf{r} \in S \equiv \{\forall \mathbf{v} \in \mathbb{R}^3 : \|\mathbf{v}\| = 1\}$, with a rotation-invariant Gaussian kernel:

$$P(\mathbf{R}; \tau, \lambda_{\parallel}, \lambda_{\perp}) = \iint_{\mathcal{S}} \Phi(\mathbf{v}) \kappa(\mathbf{R}, \mathbf{v}; \lambda_{\parallel}, \lambda_{\perp}, \tau) d\mathbf{v}, \quad (4)$$

where $\kappa(\mathbf{R}, \mathbf{v}; \lambda_{\parallel}, \lambda_{\perp}, \tau)$ has the functional form in Equation (1) for a DT, D , whose main eigenvector and eigenvalue are, respectively, $\mathbf{v}$ and $\lambda_{\parallel}$, and whose two remaining eigenvalues are both $\lambda_{\perp}$. Equation (4) is a spherical mixture model, so that $P(\mathbf{R}; \tau, D)$ can be sampled by randomly picking an orientation $\mathbf{v}$ according to the fODF, $\Phi(\mathbf{v})$, then building D from its eigenspace, and finally using the aforementioned procedure to draw a Gaussian sample. The fODF is most often represented as an expansion in the basis of real-valued, symmetric SH over the manifold S , $\tilde{Y}_l^m(\theta, \varphi)$:²²

$$\begin{aligned} \Phi(\mathbf{r}) &\equiv \tilde{\Phi}(\theta, \varphi) = \sum_{l=0}^L \sum_{m=-l}^l \phi_l^m \tilde{Y}_l^m(\theta, \varphi) \\ &\equiv \sum_{l=0}^L \sum_{m=-l}^l \phi_l^m \sqrt{\frac{2l+1}{2\pi} \frac{(l-m)!}{(l+m)!}} P_l^m(\cos(\theta)) \mathcal{T}(m\varphi), \end{aligned} \quad (5)$$

where $\mathbf{r} = [\sin(\theta)\cos(\varphi), \sin(\theta)\sin(\varphi), \cos(\theta)]^T$, $0 \leq \theta < \pi$, $0 \leq \varphi < 2\pi$ with the usual convention for spherical coordinates, P_l^m are the associated Legendre polynomials, and $\mathcal{T}$ is either the sine ($m > 0$), the cosine ($m < 0$), or $2^{-1/2}$ ($m = 0$).

2.2 | The rejection method for sampling arbitrary PDFs

Sampling of arbitrary EAPs reduces to the generation of either 3-D (Cartesian) or 2-D (spherical) random samples from known PDFs. We will use the rejection method described in,²⁴ which is illustrated in Figure 1 for a 1-D RV, $\mathbf{X}$. The method involves the target PDF of $\mathbf{X}$, $f_{\mathbf{X}}$, an

auxiliary PDF $f_{\mathbf{U}}$ of a 1-D RV $\mathbf{U}$ that we know how to sample and a constant $c > 1$ fulfilling $f_{\mathbf{X}}(t) \leq c \cdot f_{\mathbf{U}}(t)$, $\forall t$, see Figure 1A. For each instance u of $\mathbf{U}$, the dimensionality of the problem is increased with an auxiliary RV $\mathbf{V}$ which is uniformly distributed when conditioned to $\mathbf{U}$, i.e., $f_{\mathbf{V}|\mathbf{U}}(\mathbf{v}|u) = (c \cdot f_{\mathbf{U}}(u))^{-1}$, Figure 1B. The 2-D RV $\{\mathbf{U}, \mathbf{V}\}$ is uniform within $\Omega' \equiv \{\forall u, \mathbf{v} : 0 \leq \mathbf{v} \leq c \cdot f_{\mathbf{U}}(u)\}$, since $f_{\mathbf{U}, \mathbf{V}}(u, \mathbf{v}) = f_{\mathbf{V}|\mathbf{U}}(\mathbf{v}|u)f_{\mathbf{U}}(u) = (c \cdot f_{\mathbf{U}}(u))^{-1}f_{\mathbf{U}}(u) = c^{-1}$, see Figure 1C. Note both $\mathbf{U}$ and $\mathbf{V}|\mathbf{U}$ are trivially sampled (the former by hypothesis, the latter because it is uniform).

The method then rejects all samples $\{u, \mathbf{v}\}$ that lie outside $\Omega \equiv \{\forall x, y : 0 \leq y \leq f_{\mathbf{X}}(x)\}$, while those inside Ω are accepted as $\{x = u, y = v\}$, i.e., $\{\mathbf{X}, \mathbf{Y}\} \equiv \{\mathbf{U}, \mathbf{V}\} | \{\mathbf{U}, \mathbf{V}\} \in \Omega$. As Figure 1D depicts, this results in a 2-D RV $\{\mathbf{X}, \mathbf{Y}\}$ that is uniformly distributed in the new domain Ω , as long as $\Omega \subset \Omega'$ (which is guaranteed by the condition on c). It remains to discard the additional dimension $\mathbf{Y}$ to retrieve a 1-D variable $\mathbf{X}$ whose marginal PDF is $f_{\mathbf{X}}(x) = \int_0^{f_{\mathbf{X}}(x)} 1 \cdot dy$, see Figure 1E.

This same idea can be extrapolated to any number of RV, by adding 1 additional dimension, provided we can (i) evaluate the N -D PDF (e.g., using eqs. (2) or (5)), (ii) find a suitable N -D wrapper PDF that we know how to sample (say, Equation 1), and (iii) find a constant $c > 1$ such that the domain of the $(N+1)$ -D RV expanded from the latter PDF embeds the $(N+1)$ -D domain of that expanded from the former.

3 | METHODS

3.1 | Materials

For the creation of the phantom to simulate with, we had access to an *ex-vivo* pig's heart left ventricle which presented signs of ischemic cardiomyopathy scar. The acquisitions included dMRI and T1/T2 relaxometry.

Diffusion images were obtained with a Stejskal-Tanner Pulsed-Gradient SE (PGSE) sequence ($\Delta/\delta = 55.7/16.6$ ms), followed by an Echo Planar Imaging (EPI) readout. Aside from a baseline acquisition, 192 DWIs were distributed among 6 shells, namely: 300, 600, 900, 1200, 2000, and 3600 s/mm². FOV was set to 140 mm, getting 126 lines through SENSE factor 2 resulting in a voxel size of $1.09 \times 1.09 \times 1.1$ mm³. Flip angle = 90°, TE = 113.725 ms, and TR = 50,033 ms. The raw images underwent a standard preprocessing pipeline that consisted on denoising and Gibbs ringing correction through Marchenko-Pastur and PCA techniques,²⁵⁻²⁷ followed by correction of eddy currents, motion, and B0/B1 inhomogeneities.²⁸⁻³⁰

T1-relaxometry consisted in a 3D Look-Locker sequence with 39 different inversion times (TI), starting




FIGURE 1 A 1-D sketch of the rejection method. (A) f_X is the target PDF, f_U is a PDF that can be easily sampled (e.g., the Gaussian in the picture), and $c > 1$ is a constant that ensures $c \cdot f_U(t) \geq f_X(t)$, $\forall t$. (B) For any instance u generated according to f_U , a uniformly distributed sample v is generated within the range $(0, c \cdot f_U(u))$. (C) By systematically doing so, we get a uniform sampling of the $(1+1)$ -D variable $\{U, V\}$ within the domain Ω' . (D) The method rejects all samples outside the domain Ω whose upper boundary is defined by f_X , which yields a $(1+1)$ -D variable $\{X, Y\}$ uniformly distributed in Ω itself. (E) Finally, the component Y is discarded, so that the marginal variable X is distributed according to $f_X(x)$ as we pursued.

at 250 ms and separated each by 133 ms. Flip angle = 5° , TE = 2.88 ms and TR = 6.043 ms. Voxel size was $0.972 \times 0.972 \times 0.97$ mm³. T2-relaxometry used a 3D GraSE acquisition with 18 echoes split by 13 ms. TR = 1500 ms; Flip angle = 90° ; resolution: $0.97 \times 0.97 \times 0.97$ mm³. T_1 and T_2 maps were obtained from the relaxometry sequences by least-squares fitting using a custom script. Since a specific PD acquisition did not take place, the spin density map, ρ , was estimated as a function of the diffusion baseline, S , T_1 and T_2 assuming an ideal behavior:

$$S \propto \rho(1 - e^{-TR/T_1})e^{-TE/T_2} \Rightarrow \rho \propto \frac{S}{(1 - e^{-TR/T_1})e^{-TE/T_2}}. \quad (6)$$

The protocols indicated above were applied to three slices at different heights in the long axis direction.

3.2 | Computation of EAP models

From the test data described in Section 3.1, we compute per-voxel EAPs in both Cartesian and spherical coordinates as described in Section 2.1. To compute the coefficients a_{n_1, n_2, n_3} in Equation (2) we used the

Laplacian-Weighted MAP-MRI (MAPL³¹), with $N_{\max} = 6$ as the maximum order for the series expansion, a penalty term for the Laplacian computed via cross-validation, and 631 positivity constraints defined over the normalized domain of the EAP. Kernel parameters $\lambda_{\parallel}$, $\lambda_{\perp}$ in Equation (4) were computed using the procedure described in³² based on the spherical means calculated over each shell. Afterward, the coefficients ϕ_l^m up to degree $L = 6$ in Equation (5) were obtained by SH-based spherical deconvolution.²² In both cases, we used the open-source Matlab[®] implementations provided in <https://www.lpi.tel.uva.es/dmrlab>.

3.3 | Simulation of spins movements

As in Figure 1A, we need a PDF f_U that can wrap the target PDF f_X once multiplied by c . For fODF sampling, since the target PDF is defined in the compact-supported manifold S , we use a uniform PDF¹ in S or, equivalently, $f_{U,V}(\theta, \varphi) = (\sin(\theta))^{-1}$. For MAP-MRI sampling, as Figure 2A shows, the basic functions have an unbounded support at each normalized coordinate, since they are calculated as the product of a standard Gaussian PDF, $\mathcal{N}(0, 1)$, with a




FIGURE 2 Illustrative examples of random sampling of synthetic PDFs mimicking realistic anatomical configurations ((A) through (D) for MAP-MRI and (E) through (G) for SC): (A) shows some of the first (normalized) basis functions used to expand the EAP in Equation (2), together with the auxiliary (Gaussian) *wrapper* function f_U used by the rejection method at each normalized coordinate. Histograms of the samples generated according to Equation (2) are compared to the target marginal PDFs in (B) and (C), respectively, for axes ‘x’ and ‘z’ axes, and 2-D marginalization on the ‘x-z’ plane in (D). Corresponding results for Equation (5) are the marginal PDFs of each spherical coordinate in (E) and (F) and the joint target PDF in (G).

polynomial. Therefore, our *wrapper* function is designed for each normalized coordinate as the product of the $\mathcal{N}(0, 1)$ times an exponential $\exp(\kappa \cdot t^2)$, $0 < \kappa < 1/2$; this exponential grows faster than a polynomial of any degree, hence the *wrapper* function decays to 0 more slowly than the Hermite functions and it can upper-bound the target PDF in its entire domain. The product of both exponentials reads (except for constants): $g(t; \nu) = \exp(-t^2/2\nu^2)$, with $\nu = (1 - 2\kappa)^{-1/2} > 1$, so that the final wrapper function reads $g(x; \nu) \cdot g(y; \nu) \cdot g(z; \nu)$. In practice, we test discrete values of ν from 1.1 to 1.5 and keep the one leading to the smallest value of c .

Once the *wrapper* function is defined in each case, the constant c is empirically calculated as the maximum of f_X/f_U within a fixed discrete grid $([-5, 5])^3$ for MAP-MRI, as Figure 2A suggests, or $[0, \pi) \times [0, 2\pi)$ for fODFs.

For MAP-MRI, once the random sample $[x', y', z']^T$ is generated, it is rotated back by inverting Equation (3): $\mathbf{R} = \mathbf{U} \mathbf{R}'$. For SC, from the random orientation $[\theta, \phi]$, we build: $\mathbf{u}_1 = [\sin(\theta) \cos(\phi), \sin(\theta) \sin(\phi), \cos(\theta)]$, $\mathbf{u}_2 = [-\sin(\phi), \cos(\phi), 0]$, $\mathbf{u}_3 = \mathbf{u}_1 \times \mathbf{u}_2$. The DT is calculated as $D = \mathbf{U} \Lambda \mathbf{U}^T$, for $\mathbf{U} = [\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3]$, $\Lambda = \text{diag}(\lambda_{\parallel}, \lambda_{\perp}, \lambda_{\perp})$. A 3-D Gaussian sample $\mathbf{S}$ with independent components $\mathbf{S}_{\{x,y,z\}} \sim \mathcal{N}(0, 1)$ is created, and finally $\mathbf{R} = D^{1/2} \mathbf{S}$. We are

not enforcing strict non-negativity of the PDFs to sample:³³ in case negative lobes appear due to numerical issues, the rejection method will crop them by construction, since any candidate value within $[0, c \cdot g(x)]$ will be positive. Figure 2 illustrates the consistency of the method to produce random samples with both approaches.

3.4 | Simulated diffusion sequence

KomaMRI¹¹ requires both a phantom and a sequence. The phantom consists of *spins*; for each we have to provide an initial position, a value of the tuple (T_1, T_2, PD) and, if moving, a trajectory. In our case, the information provided by the EAP is the displacement at the end of the Δ -interval given by an instance of the variable $\mathbf{R}$. All the spins located at the same initial position have the same EAP, although we draw one instance of $\mathbf{R}$ per spin. Then, KomaMRI creates a *Path*² which conveys the spatial position of each spin for each time point needed in the simulation; with the information provided, the *Path* is the linear interpolation between the original and the final position of the spin (i.e., the original position plus the displacement). This is illustrated in Figure 3, lower-left corner. This




FIGURE 3 Depiction of acquisition sequence's design delivered by KomaMRI, Δ and T_E indicators were included afterward. X-axis represents time in milliseconds, Y-axis represents gradient strength in mT/m. RF pulses are represented in purple whilst gradient pulses are represented in blue, red, and green, representing x, y, and z axis encoding, respectively. This particular case represents a PGSE encoded by a gradient with direction $[0.9962627247, 0.0780837655, -0.03692572235]$ at a b-value of 3600 s/mm^2 . Diffusion motion is set to occur in the midst of both diffusion gradients, taken the effective diffusion time τ as the interval labeled as Δ . The representation in the lower-left corner illustrates that the spin movement is interpolated between the original and the final position within the Δ interval. The units in both axes do not apply for this representation.

procedure lets us depart from the *random-walk* model which, due to the Central Limit Theorem, tends to create Gaussian end-to-end displacements for a large number of steps in the walk. Finally, as initial positions we used the 2D pixel grid of each slice. Each grid position within the myocardium was populated with 2000 spins.

As for the sequence, we use PGSE with an EPI readout, see Figure 3. As the current KomaMRI version has no multicoil support, SENSE was bypassed by interlacing two different acquisitions with a blip of difference. Non-selective

excitation is used since the phantom consists of a single slice.

3.5 | Numerical validation

A three-fold quality assessment procedure was followed:

1. To check that the simulation was recreating the same orientations, a standard color-code representation was

computed. For that purpose, a DT model was calculated over each model's simulation.

2. To verify that the simulation was consistent with the source data, pixelwise Pearson correlation maps were computed. Specifically, for each shell, we define a vector per pixel consisting of the DWI values across the gradients. For each correlation value, we carry out a Fisher transformation and execute a unilateral test of equality of correlation (i.e., the correlation is at least as high as a tested value). Correlation maps were computed between the simulation and the DWI synthesized by each model, on the one hand, and between the simulation and the acquired DWI, on the other.
3. We computed the normalized mean squared error:

$$NMSE(\mathbf{x}, s) = \frac{\sum_{j=1}^{N_g^s} \left(S_s(\mathbf{x}, q_j^s) - S_{\text{ref}}(\mathbf{x}, q_j^s) \right)^2}{\sum_{j=1}^{N_g^s} S_{\text{ref}}(\mathbf{x}, q_j^s)^2} \quad (7)$$

with $\mathbf{x}$ the pixel position, s the shell index, N_g^s the number of gradients in shell s , q_j^s the j -th gradient in that shell, $S_s(\mathbf{x}, q_j^s)$ the simulated DWI at pixel $\mathbf{x}$ for gradient q_j^s and EAP of the selected model, and S_{ref} is the DWI value of the reference (for the same pixel and gradient). We used the same two references as for the correlation, i.e., either the DWI synthesized by the selected model or the acquired DWI.

To eliminate the influence of scale factors a normalization of the simulated DWIs with respect to the original acquisition was made beforehand.

4 | RESULTS

Results are presented in Figures 4 and 5. The former shows color-orientation maps following the standard RGB code for DT—red for right-left, blue for inferior-superior and green for anterior-posterior—in the two first rows and color-correlation maps in the last two. As for the orientation maps, the first row corresponds to b-value 600 and the second one to b-value 2000 s/mm². Correlation maps use a color temperature code, following blue-green-yellow-red-purple for intervals [0.5,0.6) until [0.9,1] with a step of 0.1 (see the color bars on the right-hand side). Each pixel shows a color value that corresponds to the maximum correlation accepted by the hypothesis test. Black pixels are those for which the correlation is not at least 0.5. Each column depicts results for a different diffusion model, as the legend shows.

Figure 5 depicts boxplots of both NMSE (left column) and correlation (right column). In all the cases, the variable in the horizontal axis is the b-value of each shell.

The first row shows boxplots between the simulated DWI and the DWI generated by each diffusion model while the second row compares simulation and actual DWI values. The third row shows the NMSE and the correlation for MAP-MRI with respect to the actual values parameterized by the number of spins in the simulation.

5 | DISCUSSION AND CONCLUSIONS

The color-orientation maps from Figure 4 confirm the coherence of the simulation for the three models (DTI, MAP, and SC). The correlation maps with respect to the DWIs synthesized from each model EAP (third row) show that DTI and MAP-MRI simulations closely follow their respective models while SC shows scattered areas with lower correlation. Since models are estimated from noisy data—which explains isolated model mismatches—these results mean that the simulation, in essence, does as expected. As for correlation with respect to the acquired DWIs, values are lower although without obvious visual differences between them.

The boxplots in Figure 5 show a similar pattern to that shown in Figure 4. Taking as a reference the model—subfigures (A) and (B)—the three diffusion schemes show similar behaviours, although SC shows a worse fit. In comparison with the original acquisition—subfigures (C) and (D)—MAP-MRI is consistently better than DTI and SC for both NMSE and correlation for all the shells. Hence, MAP-MRI was selected for comparing performance parameterized by the number of spins; the results in subfigures (E) and (F) show that the improvement in NMSE and correlation with the number of spins grows for higher shells. In terms of NMSE, 2000 spins seem a sensible trade-off between performance and computational complexity. We have used this number in the experiments referred to above.

In terms of computation time, EAP sampling and acquisition simulation must be differentiated since they are carried out in CPU and GPU, respectively. The number of grid positions per slice is 1757, 1739, and 711, respectively, for slices 1 through 3. The computation of 2000 samples per grid position took on the order of 7–8 ms in both an AMD Epyc 7513 and an Apple M1-Pro processor. As for simulation, each grid position took 27 ms in an NVIDIA RTX A5000. Overall, each DWI took roughly 1 min for slices 1 and 2.

Overall, we have chosen two non-Gaussian diffusion models, which make use of both Cartesian and spherical coordinates, and we have shown that both the EAP sampling and the spin-oriented simulation with a PGSE-EPI sequence provide acceptable simulation results for a wide




FIGURE 4 Color orientation maps and color correlation maps. Beginning from the left, first column represents DTI model results, second column MAP-MRI, and last SC for each row. First row represents color orientation maps for shell of b-value 600. Second row represents color orientation maps for shell of b-value 2000. Each map is obtained by adjusting a DTI model to the simulated DWIs in the indicated shell (using the EAP for the simulation corresponding to each model). The last two rows represent correlation maps against the respective model and the original acquisition, respectively, for the 2000 s/mm² shell. Color code follows color temperature, being blue a correlation coefficient of at least 0.5, green ρ of at least 0.6, yellow ρ of at least 0.7, red ρ of at least 0.8, and purple ρ of at least 0.9.

range of b-values. Moreover, this methodology shows generality so it carries over to other diffusion models and anatomical locations. As for the diffusion models themselves, MAP-MRI has shown better performance than

both DTI and SC. This leaves room for further research on the suitability of higher order diffusion models in the cardiac case. Accordingly, simulation would provide means to anticipate the result in non-ideal hardware




FIGURE 5 Boxplots comparing point-to-point metrics aggregated across three slices. Left column, subfigures (A,C,E): NMSE. Right column, subfigures (B,D,F): correlation. In all the figures, the horizontal axis represents the b-value. First row, subfigures (A,B) compare each simulation result with its corresponding model. Second row, subfigures (C,D) compare simulation with the original acquisition. Bottom row, subfigures (E,F) illustrate the metrics results for MAP-MRI and the original acquisition parameterized by the number of simulated spins—from 100 to 10 000—per grid position.

conditions, presence of *off-resonance* effects, usage of different gradient waveforms or, generally speaking, of different sequences. The limitation is that the simulation

remains valid as long as the diffusion time coincides with the one that gave rise to the EAP. This is the price to pay to avoid microstructure in the simulation.

ACKNOWLEDGMENTS

This work was supported in part by the Agencia Estatal de Investigación, under Grants RTI2018-094569-B-I00, PID2020-115339RB-I00, and PID2021-124407NB-I00 and by the company ESAOTE Ltd by grant 18IQBM. The authors express their gratitude to Dr. Filgueiras-Rama and Dr. Sánchez-González for generously sharing the diffusion images that made this research possible.

CONFLICT OF INTEREST STATEMENT

The authors declare there is no conflict of interest.

DATA AVAILABILITY STATEMENT

The methods for random sampling of the EAP as represented by MAP-MRI, and the ODF represented with SH expansions, have been coded as multi-threaded C++ code interfaced to Matlab[®] (mex functions), and are publicly available as part of the open-sourced dMRI-Lab toolbox: <http://www.lpi.tel.uva.es/dmriLab>. The particular functions that implement these methods are, respectively, `map12samples.m` and `shodf2samples.m`.

ENDNOTES

¹This is equivalent to uniformly choosing random 3-D orientations $\mathbf{r}$, which is trivially accomplished by generating 3-D Gaussian samples with independent components $\mathbf{R} \sim \mathcal{N}(0, 1)$, then computing $\mathbf{r} = \mathbf{R}/\|\mathbf{R}\|$.

²See <https://juliahealth.org/KomaMRI.jl/stable/tutorial/06-DiffusionMotion/>.

ORCID

Justino R. Rodríguez-Galván  <https://orcid.org/0000-0003-2390-5390>

Carlos Castillo-Passi  <https://orcid.org/0000-0001-6227-0477>

Pablo Irarrazaval  <https://orcid.org/0000-0002-5186-2642>

Antonio Tristán-Vega  <https://orcid.org/0000-0002-4614-2501>

Carlos Alberola-López  <https://orcid.org/0000-0003-3684-0055>

REFERENCES

- Treceño-Fernández D, Calabia-del Campo J, Bote-Lorenzo ML, Sánchez EG, de Luis-García R, Alberola-López C. A web-based educational magnetic resonance simulator: Design, implementation and testing. *J Med Syst*. 2020;44:1-11.
- Treceño-Fernández D, Calabia-del Campo J, Bote-Lorenzo ML, Gómez-Sánchez E, de Luis-García R, Alberola-López C. Integration of an intelligent tutoring system in a magnetic resonance simulator for education: Technical feasibility and user experience. *Comput Methods Prog Biomed*. 2020;195:105634.
- Torheim G, Rinck PA, Jones RA, Kvaerness J. A simulator for teaching MR image contrast behavior. *MAGMA*. 1994;2:515-522. doi:10.1007/BF01766086
- Hackländer T, Mertens H. Virtual MRI: A PC-based simulation of a clinical MR scanner. *Acad Radiol*. 2005;12:85-96.
- Jochimsen T, von Mengershausen M. ODIN - Object-oriented development interface for NMR. *J Magn Reson*. 2004;10:67-78.
- Stöcker T, Vahedipour K, Pflugfelder D, Shah N. High-performance computing MRI simulations. *Magn Reson Med*. 2010;64:186-193. doi:10.1002/mrm.22406
- Xanthis CG, Venetis IE, Aletras AH. High performance MRI simulations of motion on multi-GPU systems. *J Cardiovasc Magn Reson*. 2014;16:48. doi:10.1186/1532-429X-16-48
- Xanthis C, Venetis I, Chalkias A, Ah A. MRISIMUL: A GPU-based parallel approach to MRI simulations. *IEEE Trans Med Imaging*. 2014;33:607-617.
- Graham M, Drobnjak I, Jenkinson M, Zhang H. Quantitative assessment of the susceptibility artefact and its interaction with motion in diffusion MRI. *PLoS One*. 2017;10:12:1-25. doi:10.1371/journal.pone.0185647
- Kose R, Kose K. BlochSolver: A GPU-optimized fast 3D MRI simulator for experimentally compatible pulse sequences. *J Magn Reson*. 2017;281:51-65.
- Castillo-Passi C, Coronado R, Varela-Mattatall G, Alberola-López C, Botnar R, Irarrazaval P. KomaMRI.jl: An open-source framework for general MRI simulations with GPU acceleration. *Magn Reson Med*. 2023;90:329-342. doi:10.1002/mrm.29635
- Weine J, McGrath C, Dirix P, Buoso S, Kozerke S. CMRsim-A python package for cardiovascular MR simulations incorporating complex motion and flow. *Magn Reson Med*. 2024;91:2621-2637. doi:10.1002/mrm.30010
- Alexander D, Seunarine K, Nedjati-Gilani S, et al. Camino: Open-source diffusion-MRI reconstruction and processing. *Proceedings of the 14th Scientific Meeting of the International Society for Magnetic Resonance in Medicine (ISMRM)*. ISMRM; 2006:6-12.
- Stoeck CT, von Deuster C, Genet M, Atkinson D, Kozerke S. Second-order motion-compensated spin echo diffusion tensor imaging of the human heart. *Magn Reson Med*. 2016;75:1669-1676.
- von Deuster C, Stoeck CT, Genet M, Atkinson D, Kozerke S. Spin echo versus stimulated echo diffusion tensor imaging of the in vivo human heart. *Magn Reson Med*. 2016;76:862-872.
- Callaghan PT. *Principles of Nuclear Magnetic Resonance Microscopy*. Clarendon Press; 1991.
- Basser PJ, Pierpaoli C. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *J Magn Reson*. 1996;111:209-219.
- Novikov DS, Kiselev VG, Jespersen SN. On modeling. *Magn Reson Med*. 2018;79:3172-3193.
- Nguyen DV, Li JR, Grebenkov D, Le Bihan D. A finite elements method to solve the Bloch-Torrey equation applied to diffusion magnetic resonance imaging. *J Comput Phys*. 2014;263:283-302.
- Howard AF, Mollink J, Kleinnijenhuis M, et al. Joint modelling of diffusion MRI and microscopy. *NeuroImage*. 2019;201:116014.
- Özarslan E, Koay CG, Shepherd TM, et al. Mean Apparent Propagator (MAP) MRI: A novel diffusion imaging method for mapping tissue microstructure. *Neuroimage*. 2013;78:16-32.
- Descoteaux M, Angelino E, Fitzgibbons S, Deriche R. Regularized, fast, and robust analytical Q-ball imaging. *Magn Reson Med*. 2007;58:497-510.

23. Mekkaoui C, Reese TG, Jackowski MP, Bhat H, Sosnovik DE. Diffusion MRI in the heart. *NMR Biomed*. 2017;30:e3426.
24. Devroye L. *Non-Uniform Random Variate Generation*. Springer Science+Business Media, LLC; 1986.
25. Veraart J, Novikov DS, Christiaens D, Ades-Aron B, Sijbers J, Fieremans E. Denoising of diffusion MRI using random matrix theory. *NeuroImage*. 2016;142:394-406.
26. Cordero-Grande L, Christiaens D, Hutter J, Price AN, Hajnal JV. Complex diffusion-weighted image estimation via matrix recovery under general noise models. *NeuroImage*. 2019;200:391-404.
27. Kellner E, Dhital B, Kiselev VG, Reiser M. Gibbs-ringing artifact removal based on local subvoxel-shifts. *Magn Reson Med*. 2016;76:1574-1581.
28. Andersson JL, Sotiropoulos SN. An integrated approach to correction for off-resonance effects and subject movement in diffusion MR imaging. *NeuroImage*. 2016;125:1063-1078.
29. Zhang Y, Brady M, Smith S. Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Trans Med Imaging*. 2001;20:45-57.
30. Smith SM, Jenkinson M, Woolrich MW, et al. Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage*. 2004;23:S208-S219.
31. Fick RH, Wassermann D, Caruyer E, Deriche R. MAPL: Tissue microstructure estimation using Laplacian-regularized MAP-MRI and its application to HCP data. *NeuroImage*. 2016;134:365-385.
32. Tristán-Vega A, Aja-Fernández S. Efficient and accurate EAP imaging from multi-shell dMRI with Micro-Structure adaptive convolution kernels and dual Fourier Integral Transforms (MiS-FIT). *Neuroimage*. 2021;227:117616.
33. Haije TD, Özarslan E, Feragen A. Enforcing necessary non-negativity constraints for common diffusion MRI models using sum of squares programming. *NeuroImage*. 2020;209:116405.

How to cite this article: Rodríguez-Galván JR, Villacorta-Aylagas P, Merino-Caviedes S, et al. Sampling of non-Gaussian Ensemble Average Propagators for the simulation of diffusion magnetic resonance images. *Magn Reson Med*. 2025;1-11. doi: 10.1002/mrm.70080