Graph-Regularized Multi-Contrast Convex Quantification (GR-MCQ): A Global Optimization Framework for Microstructure Imaging

Jiyun Lim[1]

Received March 16, 2026
Accepted July 5, 2026
Electronic access August 31, 2026

Diffusion MRI (dMRI) infers fine tissue characteristics from the movement of water, but it can produce measurements that are nearly indistinguishable even when composed of two different tissues. In this unclassified setting, I developed Graph-Regularized Multi-Contrast Convex Quantification (GR-MCQ) to use relaxation-sensitive T_1/T_2 comparisons together with diffusion data. This method first represents each voxel signal as a non-negative mixture of dictionary atoms, then estimates all voxel coefficients together by applying an $\ell_{2,1}$ penalty and the anatomy-guided graph Laplacian. Since the problem is convex, every minimizer is globally optimal. I solved this problem using an accelerated proximal-gradient method that adopts the step size linked to the spectral norm. In the pilot phantom containing Gaussian noise, the relative Frobenius reconstruction error decreased, while the structural similarity increased. Subsequently, 30 independent Rician-noise experiments were conducted at each of the four signal-to-noise ratio (SNR) values. These synthetic results support feasibility and demonstrate the clearest advantage, especially at low to intermediate SNR levels.

Keywords: Diffusion MRI, Microstructure Imaging, Convex Optimization, Graph Laplacian, Multi-Contrast MRI

Introduction

A dMRI scan records the movement of water within a voxel, rather than the cell structure itself. A biophysical model is required to convert that indirect signal into tissue parameters. For example, NODDI returns the intracellular volume fraction (ICVF) and the orientation dispersion index (ODI)[1,2]. However, since different sets of microscopic parameters can predict nearly identical measurements, the effectiveness of that method may be insufficient. It was reported that, despite more than 60 diffusion measurements being conducted, two biologically valid solutions exist within a wide margin of error[3]. If the acquired data lacks distinguishing information, it cannot be restored even with a more stringent convergence setting.

Conventional pipelines usually solve the model voxel by voxel, often with a non-linear routine such as Levenberg-Marquardt. This creates several connected problems. A non-convex fit can change with its initial values, while independent fits can assign different parameters to adjacent voxels in the same tract. Running the non-linear solver across a full volume is also expensive for large datasets and time-limited workflows. The result is often visible as speckle in raw NODDI maps: ICVF or ODI can jump between neighboring voxels even when the structural image shows no tissue boundary. A Gaussian or median filter can reduce the speckle, but it cannot tell a noisy fluctuation from a real edge. The same operation that smooths a homogeneous region may also blur a boundary that matters in the final map.

AMICO adopts a convex reconstruction method[4]. By replacing nonlinear microscopic fitting with sparse linear reconstruction over a predefined dictionary, sensitivity to local minima is reduced with reduced computational complexity. However, adjacent voxels remain disconnected from each other even if they belong to the same structure. Moreover, this solution cannot express contrast not included in the predefined dictionary.

Some methods move a larger portion of the model into the signal library. For instance, microstructure fingerprint analysis constructs the dictionary through Monte Carlo diffusion simulations and calculates compartment weights by using a non-negative or sparse-normalized least squares method[5,6]. The difficult inverse model then becomes a convex search over a finite parameter grid. LiFE applies the same broad strategy at the tract level. The sparse tensor connects the candidate streamline to the predicted signal. However, since LiFE uses millions of candidates, the resulting computational complexity is very high[7,8].

The diffusion-relaxometry method applies simultaneous sampling of diffusion and relaxation to obtain the possible

[1] *Seoul Academy, South Korea*

source of missing parts[9,10]. This combination is useful since two contrasts respond to different tissue properties, resulting in improved performance[11,12]. The combined method successfully separates compartments which are difficult to distinguish by only applying diffusion and improved sensitivity to the myelin-water and free-water fractions[13,14]. By examining the healthy brain measurements, it is also shown that there exists a regional dependency between transverse relaxation time and diffusion[15]. Thus, the contrasts can be considered complementary rather than interchangeable.

Spatial information can be adopted to stabilize a fit with noise. In graph-based methods, adjacent voxels are connected with weighted edges[16,17]. The mismatches with in a specific domain leads to the high penalties. The atlas study reveals significant regional variation in microstructural markers by combining structural, relaxometry, and diffusion maps[15]. Thus, in this study, I select weights depending on the anatomical location rather than applying the same smoothing throughout the entire locations.

Recently, many studies adopt the learned estimator as an alternative. A convolutional network on a q-space sampling graph can obtain parameters quickly from undersampled data, while interpretation and transfer to a different acquisition protocol remain concerns[18]. Machine-learning alternatives include q-space deep learning[19], networks inspired by sparse reconstruction[20], and uncertainty-aware image-quality transfer[21]. Graph convolutional, Transformer, and clinically feasible deep-learning models provide additional learned estimators[18,22,23], and recent surveys place these methods in the wider diffusion MRI literature[24]. They are important empirical baselines, but they require training data and may lose accuracy when the deployment protocol differs from the training protocol.

In this paper, I study a GR-MCQ to propose the global optimization framework for microscopic imaging. In GR-MCQ, three design directions are jointly considered: include both relaxation and diffusion in the dictionary, suppress unstable atom selections with a sparsity penalty, and let neighboring voxels share information only when the structural image supports that connection. I formulate the convex optimization problem and develop the acceleration proximal-gradient algorithm to solve the problem. In doing so, multi-contrast atomic and anatomy-guided coupling are considered while preserving AMICO's global optimization properties. The experiments below ask whether its structure works as intended in the synthetically created, brain-like phantom. Since this study does not contain patient data, I limit the claim to algorithmic validation only.

Table 1 Main mathematical symbols used in the GR-MCQ formulation.

Symbol	Meaning
N	Number of voxels in the image domain
K	Number of multi-contrast measurements per voxel
M	Number of atoms in the relaxation-diffusion dictionary
$S \in R^{K\times N}$	Measured signal matrix
$D \in R^{K\times M}$	pre-defined dictionary matrix
$X \in R^{M\times N}$	non-negative coefficient matrix to be estimated
$L \in R^{N\times N}$	Anatomical graph Laplacian
λ	Joint-sparsity regularization weight
μ	Graph-Laplacian regularization weight
η	Gradient step size used by the optimization algorithm

System Model

Let each voxel $v \in \{1,\dots,N\}$ have a measured multi-contrast signal vector $s_v \in R^K$, where K denotes the total number of acquisitions. The multi-contrast signal vector s_v is modeled as a linear mixture model:

$$s_v = Dc_v + \varepsilon_v, \quad c_v \in R^M, \quad c_v \succeq 0, \tag{1}$$

where $D \in R^{K\times M}$ is a pre-defined dictionary whose columns $\{d_m\}_{m=1}^M$ represent candidate compartment atoms, c_v is the coefficient vector defined as non-negative partial-volume fractions, and ε_v represents noise and model mismatch. In this paper, both additive zero-mean Gaussian noise to isolate optimization behavior and Rician magnitude noise are adopted. The latter is more appropriate for single-coil or SENSE-combined magnitude MRI data; multi-coil sum-of-squares data may require a noncentral-χ model[25–27].

Extended multi-contrast atom design

Each dictionary atom d_m is obtained by multiplying a diffusion response by two relaxation factors. The following diffusion response is adopted to acquire k:

$$\mathrm{Diff}_{(k,m)} = \exp\left[-b_k\left(D_{\perp,m} + \left(D_{\parallel,m} - D_{\perp,m}\right)\left(g_k^\top u_m\right)^2\right)\right], \tag{2}$$

where b_k is the diffusion weighting, g_k is the unit gradient direction, and u_m represents the orientation assigned to atom m. The complete atom is modeled as

$$d_{(k,m)} = \mathrm{Diff}_{(k,m)} \exp\left(-\frac{TE_k}{T_{2,m}}\right)\left(1-2\exp\left(-\frac{TI_k}{T_{1,m}}\right)\right), \quad (3)$$

where the two added factors demonstrate inversion-recovery and echo-time dependence through T_1 and T_2. Some diffusion-similar configurations can be effectively separated due to these factors, while degeneracy is not totally eliminated. After removing physically redundant combinations, $M = 120$ atoms are retained with five orientation groups, three axial diffusivities, two radial diffusivities, two T_1 values, and two T_2 values. The grid covers $\theta \in [0,\pi)$, $\phi \in [0,2\pi)$, $D_{\parallel} \in \{1.3,1.7,2.1\} \times 10^{-3}\,\mathrm{mm}^2/\mathrm{s}$, $D_{\perp} \in \{0.3,0.7\} \times 10^{-3}\,\mathrm{mm}^2/\mathrm{s}$, $T_1 \in \{900,1300\}\,\mathrm{ms}$, and $T_2 \in \{70,100\}\,\mathrm{ms}$. For the reproducible Rician experiment, 96 measurements are generated from the Cartesian product of $b = \{0,1000,2000,3000\}\,\mathrm{s/mm}^2$, six uniformly spaced in-plane gradient directions, $TE = \{60,100\}\,\mathrm{ms}$, and $TI = \{1500,2500\}\,\mathrm{ms}$. Every dictionary column is normalized to unit l_2 norm before reconstruction. These values specify the synthetic experiment and are not intended as scanner-specific tissue constants.

Stacked global model

By stacking both the voxel signals and their coefficient vectors as $S = [s_1,\ldots,s_N] \in R^{K\times N}$ and $X = [c_1,\ldots,c_N] \in R^{M\times N}$, respectively, Equation (1) becomes one global system:

$$S = DX + E, \qquad X \succeq 0, \quad (4)$$

The columns of $E = [\varepsilon_1,\ldots,\varepsilon_N]$ collect the corresponding noise and mismatch terms.

Anatomical graph and Laplacian

A weighted voxel graph $\mathscr{G} = (\mathscr{V},\mathscr{E},W)$ is defined to model the spatial coupling from the high-resolution structural image, with $w_{ij} \geq 0$. With its degree matrix $\Delta = \mathrm{diag}\left(\sum_j w_{ij}\right)$, the combinatorial Laplacian can be modeled as

$$L = \Delta - W, \qquad L \succeq 0. \quad (5)$$

Each edge weight can be obtained as

$$w_{ij} = \exp\left(-\frac{\|r_i - r_j\|_2^2}{2\sigma_d^2}\right)\exp\left(-\frac{(I_i - I_j)^2}{2\sigma_I^2}\right), \quad j \in N_6(i), \quad (6)$$

where r_i is the coordinate of voxel i and I_i is its normalized structural intensity. The 2D phantoms use four neighbors. The notation $N_6(i)$ indicates the corresponding six-neighbor construction for a 3D volume. In this way, nearby voxels with similar structural intensity receive strong coupling. A boundary reduces the weight and permits a sharper change in the estimated coefficients.

Convex Optimization Formulation

The global coefficient matrix $X \in R^{M\times N}$ can be obtained as

$$\begin{aligned} \min_{X\in R^{M\times N}} \quad & \frac{1}{2}\|S - DX\|_F^2 + \lambda\|X\|_{2,1} + \frac{\mu}{2}\mathrm{Tr}(XLX^\top) \\ \text{s.t.} \quad & X \succeq 0, \end{aligned} \quad (7)$$

where $\lambda > 0$ and $\mu > 0$ are regularization parameters and

$$\|X\|_{2,1} \sum_{m=1}^{M} \|x^m\|_2, \qquad \|Z\|_{2,\infty} \max_m \|z^m\|_2, \quad (8)$$

with x^m denoting the m-th row of X.

In the objective function in Equation (7), the first term enforces fidelity to the observed data by penalizing the mismatch between the measured signal S and its reconstruction DX. The second term can suppress an atom across the image rather than allowing it to appear only in a few noisy voxels. The last term measures the similarity between adjacent coefficient vectors via L. In structurally homogeneous regions, the effect increases, while in areas near the edges where the corresponding weight is small, the effect weakens. By appropriately controlling the values of the weights of these three terms in Equation (7), the tradeoff among data fit, shared atom support, and spatial continuity can be adjusted.

Theorem 1 (Convexity of (7)). The objective function in (7) is convex in X. Since the constraint set $\{X : X \succeq 0\}$ is also convex, (7) constitutes a convex optimization problem.

Proof. Convexity is established term by term. (i) Data fidelity. The map

$$X \mapsto \frac{1}{2}\|S - DX\|_F^2$$

is a convex quadratic function of X. (ii) Joint sparsity. $\|X\|_{2,1}$ is a norm and is therefore convex. (iii) Graph Laplacian smoothness. Since $L \succeq 0$, there exists a matrix B such that $L = B^\top B$. Consequently,

$$\mathrm{Tr}(XLX^\top) = \|XB^\top\|_F^2,$$

which is a sum of squared norms and therefore convex. (iv) Constraint set. The set $\{X : X \succeq 0\}$ is an intersection of half-spaces and is therefore convex. Combining these four properties establishes that (7) is convex.

Remark 1 (Global Minimums and Uniqueness). Convexity makes every local minimum of (7) a global minimum. Nonetheless, the uniqueness of the optimal solution is not guaranteed. To achieve uniqueness, every nonzero feasible direction H on the active support would need

$$\|DH\|_F^2 + \mu\,\mathrm{Tr}(HLH^\top) > 0,$$

together with the relevant support conditions from the nonsmooth $l_{2,1}$ term. In this paper, a restricted minimum eigenvalue is not calculated and the active-support conditions for each experiment are also not verified. Thus, only the global optimality is claimed and not uniqueness.

Optimization Algorithm

The constrained primal problem (7) can be solved directly. The smooth part in the objective function can be given as

$$h(X) = \frac{1}{2}\|S - DX\|_F^2 + \frac{\mu}{2}\operatorname{Tr}(XLX^\top), \tag{9}$$

so that

$$\nabla h(X) = D^\top(DX - S) + \mu XL. \tag{10}$$

The remaining component is

$$\lambda\|X\|_{2,1} + I_{X\succeq 0}(X),$$

where $I_{X\succeq 0}$ is the indicator of the non-negative orthant. For each row, the proximal update is

$$\operatorname{prox}_{\eta(\lambda\|\cdot\|_{2,1}+I_+)}(y) = \begin{cases} \left(1 - \dfrac{\eta\lambda}{\|y_+\|_2}\right) y_+, & \|y_+\|_2 > \eta\lambda, \\ 0, & \|y_+\|_2 \le \eta\lambda, \end{cases} \tag{11}$$

where

$$y_+ = \max(y, 0)$$

is evaluated element by element. The positive part enforces the constraint, while the row threshold applies joint sparsity. No later projection is required. The gradient Lipschitz constant satisfies

$$L_h \le \|D\|_2^2 + \mu\|L\|_2, \tag{12}$$

so a conservative fixed step size is

$$\eta \le \frac{1}{L_h}.$$

In practice, power iteration estimates $\|D\|_2$ and $\|L\|_2$, and a backtracking step is triggered if the objective does not decrease. The experiments use standard FISTA acceleration. It leaves the minimizer unchanged but required fewer iterations in the synthetic tests. Algorithm 1 demonstrates the entire procedure to solve the problem (7).

Algorithm 1 Accelerated Projected Proximal-Gradient Solver for GR-MCQ

Input Data S, dictionary D, graph Laplacian L, parameters λ, μ, step size η

Output Coefficient matrix X

Initialize $X^{(0)} \leftarrow 0$, $Y^{(0)} \leftarrow X^{(0)}$, and $q_0 \leftarrow 1$
for $\{t = 0, \ldots, T-1\}$
  $G^{(t)} \leftarrow D^\top(DY^{(t)} - S) + \mu Y^{(t)}L$
  $Z^{(t)} \leftarrow Y^{(t)} - \eta G^{(t)}$
  for $\{m = 1, \ldots, M\}$
    $z \leftarrow \max(Z^{(t)}_{(m,:)}, 0)$
    if $\{\|z\|_2 > \eta\lambda\}$
      $X^{(t+1)}_{(m,:)} \leftarrow \left(1 - \dfrac{\eta\lambda}{\|z\|_2}\right) z$
    else
      $X^{(t+1)}_{(m,:)} \leftarrow 0$
    endif
  endfor
  $q_{t+1} \leftarrow \dfrac{1 + \sqrt{1 + 4q_t^2}}{2}$
  $Y^{(t+1)} \leftarrow X^{(t+1)} + \dfrac{q_t - 1}{q_{t+1}}\left(X^{(t+1)} - X^{(t)}\right)$
endfor
return $X^{(T)}$

Computational complexity

Each iteration requires one multiplication by D and one by $D^\top$, with cost $O(KMN)$ for dense dictionaries. The sparse graph multiplication XL costs $O(M|E|)$. The row-wise proximal step costs $O(MN)$. Thus the per-iteration cost is $O(KMN + M|E| + MN)$ and memory storage is $O(KM + MN + |E|)$ when the graph is stored sparsely. For 3D whole-brain applications, the dominant memory term is MN, so block-wise processing and sparse atom pruning are important implementation strategies.

Results

GR-MCQ was evaluated in two settings. A controlled synthetic phantom provides known coefficients for direct error measurement. A second, in-vivo-style anatomical phantom is examined with its construction labels withheld, so the analysis must rely on the kinds of proxy criteria available when ground truth is unknown.

Experimental Setup and Implementation Details

I used the optimization pipeline in the Optimization Algorithm section for every reconstruction. At each projected proximal-gradient iteration, Equation (11) imposes nonnegativity and joint sparsity. For reproducibility, Table 2 records the settings held in common across the experiments.

Dictionary and multi-contrast synthesis

For both experiments, the forward operator uses the extended relaxation-diffusion dictionary defined in (3). The original Gaussian-noise pilot uses one dictionary for both signal synthesis and reconstruction. The primary Rician comparison does the same. Two additional tests move the synthesis parameters off the reconstruction grid, allowing the dictionary to be tested as an approximation rather than an exact copy of the signal model.

Graph construction

The graph is built from the structural reference image using (6) and the connectivity in Table 2. An edge receives a large weight only when its voxels are close and have similar structural intensity. Thus two voxels inside a homogeneous region influence each other more than two voxels separated by a tissue boundary, even if both pairs are immediate grid neighbors.

Regularization parameters

Unless stated otherwise, the repeated simulations use $\lambda = 1 \times 10^{-2}$ and $\mu = 1 \times 10^{-1}$. Here λ controls row sparsity, or shared atom support across voxels, and mainly affects noise robustness. The parameter μ sets the strength of graph-guided coupling and mainly affects anatomical coherence.

Synthetic Phantom Validation

Experimental Setup

A 2D synthetic phantom was built with three kinds of regions: anisotropic white-matter-like tissue, isotropic tissue water, and free-water-like compartments. Multi-contrast measurements were synthesized according to the linear model

$$S = DX_{\mathrm{GT}} + E, \tag{13}$$

where D is the relaxation-diffusion dictionary, X_{GT} is the ground-truth coefficient matrix, and E is additive zero-mean Gaussian noise. Three reconstructions were compared: voxel-wise fitting, voxel-wise fitting followed by smoothing, and GR-MCQ.

Quantitative Evaluation

Coefficient recovery accuracy is measured using the relative Frobenius reconstruction error

$$\mathrm{RFRE} = \frac{\|X_{\mathrm{est}} - X_{\mathrm{GT}}\|_F}{\|X_{\mathrm{GT}}\|_F}, \tag{14}$$

which was previously labeled NRMSE; the quantity is a normalized matrix error, computed over the full coefficient matrix, and is distinct from a root-mean-square voxel error. The structural similarity index (SSIM) is also computed between the reconstructed and ground-truth coefficient maps.

Table 3 shows GR-MCQ with the lowest RFRE and the highest SSIM among the compared methods in this pilot realization. The $l_{2,1}$ penalty suppresses noise-driven spurious atom activations, which stabilizes reconstruction, while the graph Laplacian enforces within-region consistency without sacrificing boundary sharpness where graph weights diminish. Because only a single realization is reported here, statistical significance is not claimed.

Qualitative Comparison

I inspected Figures 1–3 on the same coefficient scale. The boundary for Atom 2 looks slightly sharper than those for Atoms 1 and 3, but I did not test this visual difference for significance and therefore combine the atoms in the analysis below. The independently fitted voxel maps contain isolated errors. Smoothing removes part of that speckle, but the tissue transitions also become less distinct. In the GR-MCQ panels, many isolated fluctuations disappear while the region boundaries remain visible. That is the specific visual behavior the anatomy-weighted graph was designed to produce.

Rician-Noise Robustness and Repeated Simulations

Experimental Design

The Rician analysis used a 5×9 three-region phantom and the 120-atom dictionary. The acquisition followed the 96-measurement protocol in Table 2. For each SNR in $\{10, 20, 30, 40\}$, 30 independent magnitude-noise realizations were generated with fixed master seed:

$$S_{\mathrm{Ric}} = \sqrt{(S_{\mathrm{clean}} + N_1)^2 + N_2^2}, \qquad N_1, N_2 \overset{\mathrm{iid}}{\sim} \mathcal{N}(0, \sigma^2), \tag{15}$$

where

$$\sigma = \frac{s_{95}}{\mathrm{SNR}}$$

and s_{95} is the 95th percentile of the noise-free signal. Because atoms that differ only slightly in T_1, T_2, or diffusivity can be nearly indistinguishable, the primary Rician endpoint is computed after summing the 24 parameter variants within each of the five orientation-compartment groups:

Table 2 Parameter values used to reproduce the reported reconstructions.

Setting	Choice used in this study
Dictionary size M	120 atoms
Rician-test measurements K	96
Diffusion weightings	$b = \{0, 1000, 2000, 3000\}$ s/mm^2
Gradient directions	6 uniformly spaced in-plane directions
Relaxation contrasts	$TE = \{60, 100\}$ ms; $TI = \{1500, 2500\}$ ms
Orientation grid	5 orientation groups
Diffusivity grid	$D_{\parallel} = \{1.3, 1.7, 2.1\} \times 10^{-3}$; $D_{\perp} = \{0.3, 0.7\} \times 10^{-3}$ mm^2/s
T_1/T_2 grid	$T_1 = \{900, 1300\}$ ms; $T_2 = \{70, 100\}$ ms
Regularization λ	1×10^{-2}
Regularization μ	1×10^{-1}
Graph connectivity	4-neighbor (2D); 6-neighbor extension (3D)
σ_d (spatial)	1 voxel
σ_I (intensity)	0.15 after structural-intensity normalization
Optimization iterations	300
Step size η	$0.95/(\|D\|_2^2 + \mu\|L\|_2)$
Stopping check	300 iterations; objective recorded at each iteration
Initialization	$X^{(0)} = 0$

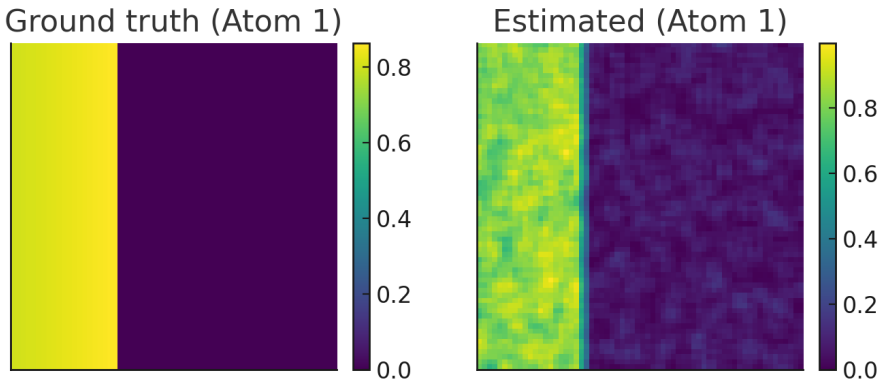

Figure. 1 Atom 1 in the synthetic phantom. Every panel uses the same coefficient range and color bar.

Table 3 Pilot quantitative comparison on one Gaussian-noise phantom realization.

Method	RFRE (%)	SSIM
Voxel-wise fitting	15.2	0.92
Voxel-wise + smoothing	13.8	0.94
Proposed (GR-MCQ)	10.5	0.98

$$\tilde{X}_{(g,:)} = \sum_{m \in G_g} X_{(m,:)}. \tag{16}$$

The group-level RFRE applies (14) to $\tilde{X}$. This endpoint tests recovery of the spatial compartment maps without claiming exact identification of every finely discretized atom.

Repeated-Noise Results

The most significant difference appears at SNR 10 which can be considered relatively noisy environment. Compared with voxel-wise fitting, the GR-MCQ reduces the paired RFRE by 3.27%, with a 95% confidence interval between 3.04 and 3.50. At SNR 20, the reduction is smaller, at 0.74% (0.61 to 0.86). At SNR 30, a difference of 0.06% is negligible, and that interval lies between -0.04 and 0.17. The tendency is reversed at SNR 40, where GR-MCQ has an error 0.56% higher. These results can be interpreted as a bias-variance trade-off. Coupling is useful when the noise is large, while regularization bias becomes significant with low noise-level. The post-smoothing results perform poorly at all SNR values since they average coefficients across the boundaries of the two sharp regions of the phantom.

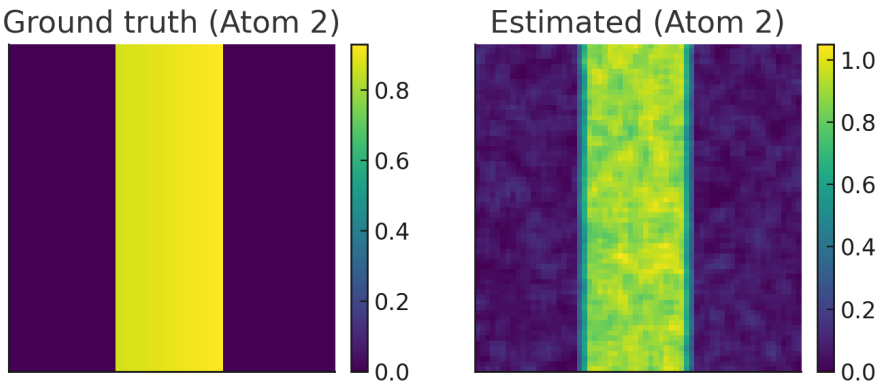

Figure. 2 Atom 2 in the synthetic phantom, shown on the coefficient scale used for Atoms 1 and 3.

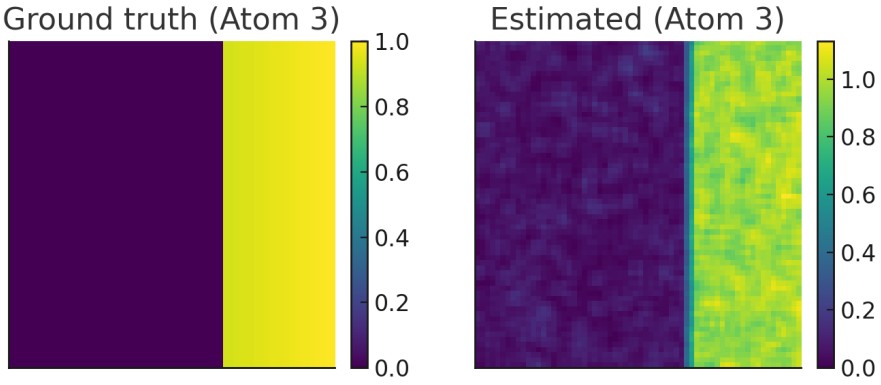

Figure. 3 Atom 3 in the synthetic phantom. The shared color scale allows a panel-to-panel comparison.

Dictionary Mismatch

To reduce inverse-crime bias, additional SNR-20 experiments synthesize signals using shifted physical parameters while reconstruction retains the original grid. The mild mismatch scales both diffusivities by 1.025 and shifts T_1/T_2 by $+25/+2.5$ ms; the moderate mismatch uses 1.05 and $+50/+5$ ms.

Table 4 GR-MCQ sensitivity to synthesis-reconstruction dictionary mismatch at SNR 20 (10 independent Rician-noise realizations).

Dictionary condition	Group RFRE (%)	SSIM
Matched	9.21 ± 0.60	0.9930 ± 0.0008
Mild mismatch	10.08 ± 0.67	0.9916 ± 0.0010
Moderate mismatch	10.92 ± 0.71	0.9902 ± 0.0012

The moderate mismatch raises group RFRE by 1.71%. The increase is clear, although the reconstruction does not fail. Since this result comes from one controlled phantom sweep, it does not establish sensitivity across other tissue models, acquisition protocols, or dictionary resolutions.

Regularization Trade-off

The graph weight μ was swept over $\{0, 0.02, 0.05, 0.10, 0.20\}$ at SNR 20 using 10 repeated noise realizations while holding $\lambda = 0.01$ fixed.

Increasing μ from 0 to 0.20 lowers mean group RFRE from 9.90% to 9.12%, while the residual ratio increases from 0.1142 to 0.1153. The curve gives a direct, quantitative view of the bias-variance trade-off, supplementing the qualitative explanation given above. The reported setting $\mu = 0.10$ sits near the bend of the curve, which avoids simply selecting the largest tested regularization value from this small phantom.

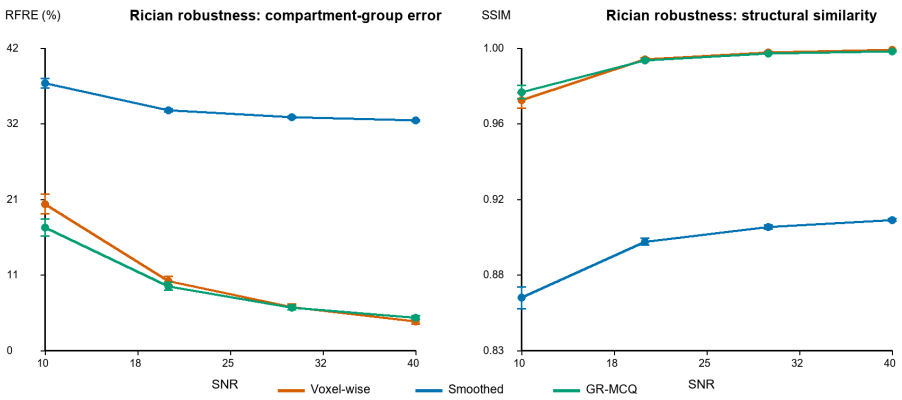

Figure. 4 Rician-noise robustness over 30 independent realizations per SNR. Error bars show one standard deviation. RFRE is evaluated on the five aggregated orientation-compartment maps defined in (16); SSIM is averaged across the three active maps.

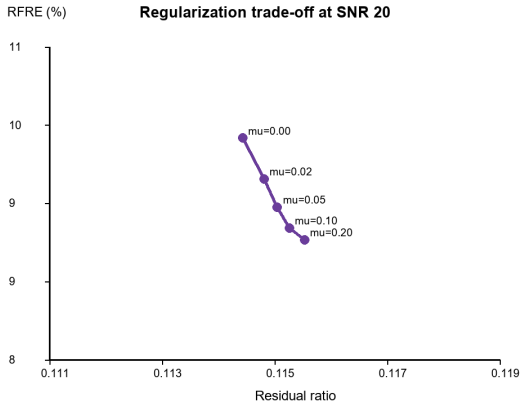

Figure. 5 Data-fidelity and reconstruction-error trade-off at SNR 20.

In-Vivo-Style Anatomical Validation

Experimental Design

A simple three-region geometry was replaced with a brain-like phantom with white matter (WM), gray matter (GM), and cerebrospinal fluid (CSF). Multi-contrast measurements are generated from tissue-specific diffusion and relaxation parameters. Bilateral graph weights are also derived from structural T1-weighted images. While analyzing, I deliberately hid the structure label. By visualizing them, residual ratio and coherence testing will become possible unlike patient scans without an actual microstructure map.

Representative Results

In Figure 6, the voxel-wise maps contain scattered coefficient changes with no corresponding feature in the structural reference. They are most noticeable around the WM-GM and GM-CSF interfaces. After graph regularization, most isolated values are gone and each tissue region is more continuous, yet the interfaces can still be seen. Equation (6) explains this result: a structural-intensity change lowers the edge weight, so coefficient vectors across the boundary pull less strongly on one another. For this phantom, the graph therefore improves within-region consistency without the broad edge blurring produced by an ordinary spatial filter.

Convergence Behavior

In Figure 7, the primal objective drops quickly at first and then levels off. With $T = 300$ iterations and the step-size rule in (12), the representative run reaches a stable plateau.

Proxy Quantitative Analysis Without Ground Truth

True in vivo microstructure parameters cannot be observed directly, so the anatomical-phantom analysis uses proxy criteria. One proxy is the relative residual fitting error,

$$\text{Residual Ratio} = \frac{\|S - DX\|_F}{\|S\|_F}, \tag{17}$$

which measures data fidelity without requiring the unknown ground truth.

Table 5 Residual fitting error in the in-vivo-style experiment.

Method	Residual Ratio
Voxel-wise fitting	0.11
Proposed (GR-MCQ)	0.12

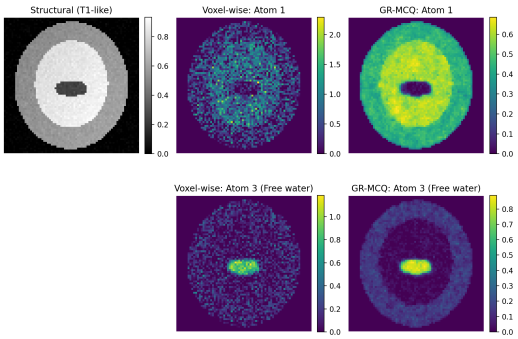

Figure. 6 Brain-like synthetic experiment. The top row contains the structural reference and voxel-wise estimates; the lower row contains the GR-MCQ result.

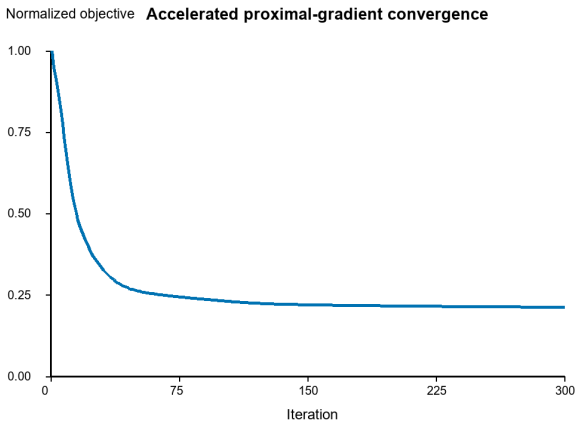

Figure. 7 Primal-objective convergence during accelerated proximal-gradient iterations. The vertical axis is normalized by the initial objective value.

Table 6 shows similar data fidelity for GR-MCQ and voxel-wise fitting. The residual ratio rises modestly, from 0.11 to 0.12. In return, the Laplacian term produces a more spatially coherent estimate. Figure 5 measures this trade-off across the tested values of μ.

Reproducibility Analysis

For each of 30 paired trials, two independent SNR-20 Rician observations were generated and reconstructed separately. Consistency of the aggregated compartment maps is measured with ICC(3,1), a two-way mixed-effects, single-measure intraclass correlation suited to this repeated-measurement comparison[28].

$$\mathrm{ICC}(3,1) = \frac{MS_R - MS_E}{MS_R + (k-1)MS_E}, \tag{18}$$

where MS_R is the between-target mean square, MS_E is the residual mean square, and $k = 2$ is the number of repeated reconstructions. This is a simulation-noise reproducibility measure, not clinical scan-rescan reliability.

Table 6 Reproducibility across 30 pairs of independent SNR-20 Rician-noise realizations.

Method	ICC(3,1) (mean $\pm$ std)
Voxel-wise fitting	0.9945 ± 0.0006
Proposed (GR-MCQ)	0.9983 ± 0.0003

Both methods are highly consistent in this simple phantom (Table 7), and GR-MCQ adds a small increase. The joint sparsity prior discourages atoms from switching between reconstructions, while graph regularization shares stable support across anatomical neighborhoods. Both effects reduce sensitivity to the particular noise realization. These simulated pairs are not a substitute for clinical scan-rescan data.

Summary of Findings

I do not read these experiments as evidence of a uniform advantage. GR-MCQ has lower error in the Gaussian pilot and in the noisier Rician trials, but the difference nearly disappears at SNR 30 and changes direction at SNR 40. In the anatomical phantom, the visible boundaries remain, and the method stays stable in the particular dictionary-mismatch and repeated-noise tests performed here. The residual increase and convergence curve are also consistent with the expected cost of regularization. On that basis, I consider the algorithm feasible for further testing. I cannot draw a clinical conclusion because this study includes neither patients nor healthy volunteers.

Discussion

The main lesson I take from the phantom experiments is that contrast information and spatial information can be added to the same convex reconstruction without forcing every neighboring voxel to agree. In GR-MCQ, the relaxation-diffusion dictionary addresses compartment ambiguity, the $l_{2,1}$ term suppresses unstable atom choices, and the structural-MRI graph couples estimates only where the anatomy gives a reason to do so.

Since the formulated optimization problem in this paper is convex, it can be solved efficiently compared with other approaches such as voxel-wise non-linear fit. Specifically, GR-MCQ reduces coefficient errors at Gaussian pilot and low-to-mid-level Rician SNR, while improving map-to-map coherency without noticeable boundary loss. Furthermore, the residual ratio increases only slightly and the consistency of repeated noise improves.

I have kept the conclusions restricted to controlled synthetic and phantom data because no patient or healthy-volunteer scans were analyzed. A clinical study would need multi-contrast human data, protocol-variation and scan-rescan tests, appropriate preprocessing, and the required ethics and consent procedures. The SNR-40 result is especially important to this limit: GR-MCQ has slightly higher group RFRE than voxel-wise fitting, so regularization can do more harm than good when the measurements are already clean. I also observed a modest error increase under the moderate dictionary mismatch. That experiment covers only two off-grid shifts and should be expanded before drawing a general conclusion about mismatch.

Three-dimensional computation remains unresolved. With a dense dictionary, one iteration costs $O(KMN + M|E|)$ and storing $X \in R^{M\times N}$ dominates memory. A whole-brain implementation will therefore need sparse dictionaries, block-wise reconstruction, GPU acceleration, or some combination of these. I did not make a matched wall-clock comparison with Levenberg-Marquardt fitting, AMICO, ADMM, or learned estimators. The complexity expression cannot substitute for timing all methods on the same hardware with the same dictionary, voxel count, and stopping tolerance. Automatic selection of (λ, μ) would also be needed for practical use. For the next study, I would first improve the dictionary rather than add another regularizer. Better microstructural priors could address the original degeneracy while retaining the convexity and interpretability that motivated this approach.

References

1 H. Zhang, T. Schneider, C. A. Wheeler-Kingshott, D. C. Alexander. NODDI: practical in vivo neurite orientation dispersion and density imaging of the human brain. NeuroImage. Vol. 61, pp. 1000–1016, 2012. https://doi.org/10.1016/j.neuroimage.2012.03.072.

2 D. S. Novikov, E. Fieremans, S. N. Jespersen, V. G. Kiselev. Quantifying brain microstructure with diffusion MRI: theory and parameter estimation. NMR in Biomedicine. Vol. 32, e3998, 2019. https://doi.org/10.1002/nbm.3998.

3 I. O. Jelescu, J. Veraart, E. Fieremans, D. S. Novikov. Degeneracy in model parameter estimation for multi-compartmental diffusion in neuronal tissue. NMR in Biomedicine. Vol. 29, pp. 33–47, 2016. https://doi.org/10.1002/nbm.3450.

4 A. Daducci, E. J. Canales-Rodríguez, H. Zhang, T. B. Dyrby, D. C. Alexander, J.-P. Thiran. Accelerated microstructure imaging via convex optimization (AMICO) from diffusion MRI data. NeuroImage. Vol. 105, pp. 32–44, 2015. https://doi.org/10.1016/j.neuroimage.2014.10.026.

5 G. Rensonnet, B. Scherrer, G. Girard, A. Jankovski, S. K. Warfield, B. Macq, J.-P. Thiran, M. Taquet. Towards microstructure fingerprinting: estimation of tissue properties from a dictionary of Monte Carlo diffusion MRI simulations. NeuroImage. Vol. 184, pp. 964–980, 2019. https://doi.org/10.1016/j.neuroimage.2018.09.076.

6 E. Özarslan, C. G. Koay, T. M. Shepherd, M. E. Komlosh, M. O. İrfanoğlu, C. Pierpaoli, P. J. Basser. Mean apparent propagator (MAP) MRI: a novel diffusion imaging method for mapping tissue microstructure. NeuroImage. Vol. 78, pp. 16–32, 2013. https://doi.org/10.1016/j.neuroimage.2013.04.016.

7 F. Pestilli, J. D. Yeatman, A. Rokem, K. N. Kay, B. A. Wandell. Evaluation and statistical inference for human connectomes. Nature Methods. Vol. 11, pp. 1058–1063, 2014. https://doi.org/10.1038/nmeth.3098.

8 C. F. Caiafa, F. Pestilli. Multidimensional encoding of brain connectomes. Scientific Reports. Vol. 7, 11491, 2017. https://doi.org/10.1038/s41598-017-09250-w.

9 Y. Wu, X. Liu, X. Zhang, K. M. Huynh, S. Ahmad, P.-T. Yap. Relaxation-diffusion spectrum imaging for probing tissue microarchitecture. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2023. Lecture Notes in Computer Science, Vol. 14227, pp. 152–162, 2023. https://doi.org/10.1007/978-3-031-43993-3_15.

10 J. Hutter, P. J. Slator, D. Christiaens, R. P. A. G. Teixeira, T. Roberts, L. Jackson, A. N. Price, S. Malik, J. V. Hajnal. Integrated and efficient diffusion-relaxometry using ZEBRA. Scientific Reports. Vol. 8, 15138, 2018. https://doi.org/10.1038/s41598-018-33463-2.

11 P. J. Slator, M. Palombo, K. L. Miller, C.-F. Westin, F. Laun, D. Kim, J. P. Haldar, D. Benjamini, G. Lemberskiy, J. P. de Almeida Martins, J. Hutter. Combined diffusion-relaxometry microstructure imaging: current status and future prospects. Magnetic Resonance in Medicine. Vol. 86, pp. 2987–3011, 2021. https://doi.org/10.1002/mrm.28963.

12 S. Coelho, Y. Liao, F. Szczepankiewicz, J. Veraart, S. Chung, Y. W. Lui, D. S. Novikov, E. Fieremans. Assessment of precision and accuracy of brain white matter microstructure using combined diffusion MRI and relaxometry. Human Brain Mapping. Vol. 45, e26725, 2024. https://doi.org/10.1002/hbm.26725.

13 S. Endt, M. Engel, E. Naldi, R. Assereto, M. Molendowska, L. Mueller, C. Mayrink Verdun, C. M. Pirkl, M. Palombo, D. K. Jones, M. I. Menzel. In vivo myelin water quantification using diffusion-relaxation correlation MRI: a comparison of 1D and 2D methods. Applied Magnetic Resonance. Vol. 54, pp. 1571–1588, 2023. https://doi.org/10.1007/s00723-023-01584-1.

14 C.-F. Westin, H. Knutsson, O. Pasternak, F. Szczepankiewicz, E. Özarslan, D. van Westen, C. Mattisson, M. Bogren, L. J. O'Donnell, M. Kubicki, D. Topgaard, M. Nilsson. Q-space trajectory imaging for multidimensional diffusion MRI of the human brain. NeuroImage. Vol. 135, pp. 345–362, 2016. https://doi.org/10.1016/j.neuroimage.2016.02.039.

15 I. S. Walimuni, K. M. Hasan. Atlas-based investigation of human brain tissue microstructural spatial heterogeneity and interplay between trans-

verse relaxation time and radial diffusivity. NeuroImage. Vol. 57, pp. 1402–1410, 2011. https://doi.org/10.1016/j.neuroimage.2011.05.063.

16 F. Zhang, E. R. Hancock. Riemannian graph diffusion for DT-MRI regularization. In: Medical Image Computing and Computer-Assisted Intervention – MICCAI 2006. Lecture Notes in Computer Science, Vol. 4191, pp. 234–242, 2006. https://doi.org/10.1007/11866763_29.,.

17 L. M. Harrison, W. D. Penny, J. Ashburner, N. J. Trujillo-Barreto, K. J. Friston. Diffusion-based spatial priors for imaging. NeuroImage. Vol. 38, pp. 677–695, 2007. https://doi.org/10.1016/j.neuroimage.2007.07.032.

18 G. Chen, Y. Hong, Y. Zhang, J. Kim, K. M. Huynh, J. Ma, W. Lin, D. Shen, P.-T. Yap. Estimating tissue microstructure with undersampled diffusion data via graph convolutional neural networks. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2020. Lecture Notes in Computer Science, pp. 280–290, 2020. https://doi.org/10.1007/978-3-030-59728-3_28.

19 V. Golkov, A. Dosovitskiy, J. I. Sperl, M. I. Menzel, M. Czisch, P. Sämann, T. Brox, D. Cremers. q-Space deep learning: twelve-fold shorter and model-free diffusion MRI scans. IEEE Transactions on Medical Imaging. Vol. 35, pp. 1344–1351, 2016. https://doi.org/10.1109/TMI.2016.2551324.

20 C. Ye. Estimation of tissue microstructure using a deep network inspired by a sparse reconstruction framework. In: Information Processing in Medical Imaging. Lecture Notes in Computer Science, pp. 466–477, 2017. https://doi.org/10.1007/978-3-319-59050-9_37.

21 R. Tanno, D. E. Worrall, A. Ghosh, E. Kaden, S. N. Sotiropoulos, A. Criminisi, D. C. Alexander. Bayesian image quality transfer with CNNs: exploring uncertainty in dMRI super-resolution. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2017. Lecture Notes in Computer Science, pp. 611–619, 2017. https://doi.org/10.1007/978-3-319-66182-7_70.

22 T. Zheng, G. Yan, H. Li, W. Zheng, W. Shi, Y. Zhang, C. Ye, D. Wu. A microstructure estimation Transformer inspired by sparse representation for diffusion MRI. Medical Image Analysis. Vol. 86, 102788, 2023. https://doi.org/10.1016/j.media.2023.102788.

23 Y. Li, Z. Zhuo, C. Liu, Y. Duan, Y. Shi, T. Wang, R. Li, Y. Wang, J. Jiang, J. Xu, D. Tian, X. Zhang, F. Shi, X. Zhang, A. Carass, F. Barkhof, J. L. Prince, C. Ye, Y. Liu. Deep learning enables accurate brain tissue microstructure analysis based on clinically feasible diffusion magnetic resonance imaging. NeuroImage. Vol. 300, 120858, 2024. https://doi.org/10.1016/j.neuroimage.2024.120858.

24 D. Karimi, S. K. Warfield. Diffusion MRI with machine learning. Imaging Neuroscience. Vol. 2, 2024. https://doi.org/10.1162/imag_a_00353.

25 H. Gudbjartsson, S. Patz. The Rician distribution of noisy MRI data. Magnetic Resonance in Medicine. Vol. 34, pp. 910–914, 1995. https://doi.org/10.1002/mrm.1910340618.

26 C. G. Koay, P. J. Basser. Analytically exact correction scheme for signal extraction from noisy magnitude MR signals. Journal of Magnetic Resonance. Vol. 179, pp. 317–322, 2006. https://doi.org/10.1016/j.jmr.2006.01.016.

27 J. Veraart, E. Fieremans, D. S. Novikov. Diffusion MRI noise mapping using random matrix theory. Magnetic Resonance in Medicine. Vol. 76, pp. 1582–1593, 2016. https://doi.org/10.1002/mrm.26059.

28 T. K. Koo, M. Y. Li. A guideline of selecting and reporting intraclass correlation coefficients for reliability research. Journal of Chiropractic Medicine. Vol. 15, pp. 155–163, 2016. https://doi.org/10.1016/j.jcm.2016.02.012.