Taming Gradient Perturbations: Probabilistic Reparameterization for Robust Physics-Informed MRI Protocol Optimization

Beomgu Kang [1] **Hye-Young Heo** [2] **Hyunseok Seo** [1]

Abstract

Reliable estimation of tissue characteristics in quantitative MRI depends on optimizing acquisition protocols defined by combinations of control parameters. Existing physics-informed optimization methods rely on gradient-based updates, but are highly sensitive to stochastic noise arising from both the imaging system and stochastic gradient descent (SGD), often leading to unstable updates and suboptimal solutions in combinatorial control parameter spaces. To address this, we present Robust Physics-Informed Optimization (RPO), a noise-robust framework that replaces direct optimization of control parameter values with probabilistic optimization over candidate control parameters. Specifically, RPO introduces a probabilistic reparameterization of the control parameters via the Gumbel-Softmax mapping, thereby changing how gradient perturbations propagate through the Jacobian of the mapping. Theoretically, this reparameterization attenuates gradient perturbations, yielding an operator-norm bound proportional to $1/\tau$ and providing a mechanistic explanation for robustness to system noise and SGD-induced stochasticity. Furthermore, the temperature parameter τ shapes the optimization dynamics by balancing broad exploration of candidate configurations in the early iterations with progressively sharper focus on high-probability solutions. Experiments on both simulated and *in vivo* MRI data demonstrate that RPO consistently improves optimization robustness and tissue-estimation accuracy under realistic noise conditions, with successful translation to actual MRI systems.

[1]Department of Artificial Intelligence, Korea University, Seoul, Republic of Korea [2]Department of Radiology, Johns Hopkins University, Baltimore, Maryland, USA. Correspondence to: Hyunseok Seo <seoh@korea.ac.kr>.

Proceedings of the AI4Physics Workshop at the 43rd International Conference on Machine Learning (AI4Physics@ICML 2026), Seoul, South Korea. 2026. Copyright 2026 by the author(s).

1. Introduction

Optimizing the acquisition protocols of control (scan) parameters is crucial for accurately determining tissue characteristics in quantitative magnetic resonance imaging (MRI). Basically, the MRI system produces its signal by configuring a combination of control parameters. Repeated acquisitions using varied combinations of control parameters yield evolving signal trajectories, from which tissue characteristics are inferred based on the knowledge of the MRI system (Fig. 1). Since the observed signal evolutions depend on varied combinations of control parameters, optimizing control parameters is a critical research focus in MRI for accurately estimating unknown *in vivo* tissue characteristics.

Various quantitative MRI techniques are designed using their specific physics models, e.g., intravoxel incoherent motion, arterial spin labeling, and saturation transfer MRI (Van Zijl & Yadav, 2011; Le Bihan, 2019; Hilbert et al., 2020; Hernandez-Garcia et al., 2022). Although the physics model bridges tissue characteristics and signal evolutions, optimizing control parameters is challenging due to their tight coupling within the complex physics model. Since scans using different control parameter configurations can generate identical signal responses, this makes the acquisition-design objective highly non-convex and potentially multi-solution (Cohen et al., 2025). This resembles multi-solution combinatorial optimization problems, where global optimality is difficult to guarantee and practical methods often rely on heuristic or locally optimized strategies (Huang et al., 2018; 2019; Jian & Hsieh, 2022).

Recently, neural networks have been proposed to solve nonlinear inverse problems of estimating tissue characteristics (Yoon et al., 2018; Scannell et al., 2020; Jung et al., 2022; Kang et al., 2021). Therefore, the tissue estimation error can be computed by a neural network using simulated signals generated from the physics-based signal model, allowing the control parameter configurations to be updated via gradient descent in a way that minimizes the estimation error of tissue characteristics (Lahiri et al., 2020; Lee et al., 2021; Planchuelo-Gómez et al., 2024; Cohen et al., 2025). However, the standard gradient descent method has difficulty in distinguishing the signal from noise (Li et al., 2020; Song et al., 2022; Meng et al., 2025), making it prone to system

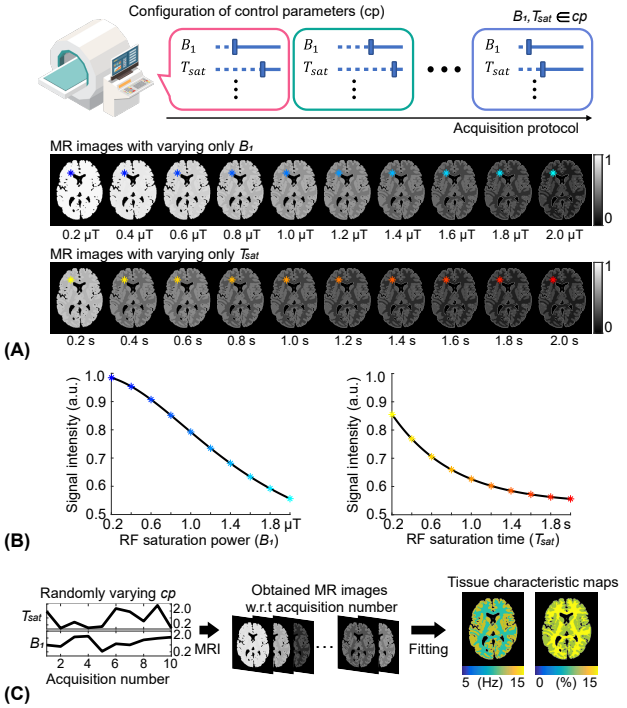

Figure 1. An overview of the quantitative MRI technique. **(A)** Two series of MR images are obtained by varying RF saturation power (B_1) and saturation time (T_{sat}), both of which modulate the signal intensity. **(B)** Voxel-wise signals (stars) are plotted against a varying control parameter, with solid curves representing the physics-based fits. **(C)** Data from multiple acquisitions with distinct control parameter configurations (cp) are fitted to the physics-informed signal model to estimate tissue maps. The BrainWeb digital phantom (Collins et al., 1998) was utilized for this visualization.

noise and also to the inherent stochasticity of the stochastic gradient descent (SGD) process (Li et al., 2017; Smith & Le, 2017; Evans et al., 2020; Smith et al., 2021). Moreover, conventional optimizations update the control parameters to specific values in the current iteration, and in the subsequent iteration, only the vicinity of these values is explored. Thus, it depends heavily on initialization and current values, leading to an inefficient search and potentially trapped in local optima.

To address these limitations, we propose a Robust Physics-informed Optimization (RPO) framework that modifies how gradient perturbations propagate during optimization. Specifically, RPO introduces a probabilistic reparameterization of the control parameters via the Gumbel-Softmax mapping. Unlike conventional approaches that directly update control parameter values, RPO updates logits that define a distribution over candidate control parameters, so that stochastic perturbations are propagated through the Jacobian of the reparameterization. This attenuates perturbations arising from both system noise and SGD, thereby improving optimization robustness. Moreover, maintaining a distribution over the combinatorial control parameter space enables broader exploration of diverse candidate configurations while progressively concentrating on high-probability solutions. We validate RPO on both simulated and real saturation transfer (ST) MRI data, where the ST model serves as the underlying physics-based signal model. ST MRI enables quantitative assessment of tissue metabolites as well as semisolid macromolecules in the brain, but requires extensive repetitive acquisitions (Perlman et al., 2023; Vladimirov et al., 2025; Singh et al., 2025), making it a challenging yet suitable setting for protocol optimization.

The main contributions of our work are summarized as follows:

- **Perturbation-Aware Robust Optimization:** Our theoretical analysis shows that RPO attenuates gradient perturbations arising from system noise and SGD-induced stochasticity through the Jacobian of the Gumbel-Softmax reparameterization. This yields an operator-norm bound proportional to $1/\tau$, which is consistently supported by our empirical findings.
- **Distribution-Based Optimization over Combinatorial Control Parameter Space:** Instead of directly updating control parameter values, RPO maintains and updates a distribution over candidate control parameters in the combinatorial control space, enabling broader exploration of diverse candidate configurations while progressively concentrating on high-probability solutions through temperature annealing.
- **Improved Optimization Performance:** RPO improves estimation accuracy for tissue characteristics and overall optimization robustness compared to existing acquisition strategies.
- **Translation to Real-World MRI Systems:** Acquisition schedules optimized via physics-based simulations successfully translate to actual MRI systems, enabling accurate quantitative tissue mapping with 6-fold acquisition acceleration.

2. Related Work

Recent works have proposed novel methods to optimize the acquisition protocols of control parameters in quantitative MRI. The correlation of signal evolutions between different tissue types was reduced to enhance the differentiation of signals corresponding to distinct tissue characteristics (Cohen & Rosen, 2017). The variance of estimates was minimized for each tissue characteristic using the Cramer-Rao lower bound (CRLB), thereby maximizing the SNR efficiency of tissue estimation (Zhao et al., 2018). A pseudo-randomly varying acquisition protocol was designed by increasing temporal incoherence and thus decreasing the in-

formation redundancy in repeated, serial acquisitions (Kim et al., 2020). However, the acquisition protocols optimized using the aforementioned indirect metrics cannot guarantee low estimation errors for tissue characteristics. Although the CRLB approach somewhat ensures high accuracy in estimating tissue characteristics, it requires calculating the derivative of the signal model with respect to each tissue characteristic, which is often challenging for complex signal models.

Learning-based optimization of acquisition schedules (LOAS) was recently proposed as a physics-informed approach for optimizing MRI protocols, directly minimizing the estimation error of tissue characteristics (Kang et al., 2022; 2026). A neural network was adopted to solve the inverse problem of complex Bloch equations for tissue characteristics estimation, and the estimation loss was backpropagated to both the control parameters and the network. This facilitated end-to-end training of both the acquisition protocol optimization and the downstream tissue characteristic estimation network.

The Gumbel-Softmax reparameterization trick has been used in prior MRI studies to optimize k-space subsampling patterns (Van Gorp et al., 2021; Kang & Seo, 2026). However, its role in our framework is fundamentally different. Prior work uses Gumbel-Softmax to parameterize binary masks for k-space subsampling, aiming to improve sampling efficiency within a single scan, whereas RPO uses it to parameterize a distribution over MR control parameters for acquisition protocol optimization across repeated scans. Importantly, in our setting, this reparameterization changes how gradient perturbations propagate during physics-informed optimization by passing them through the Jacobian of the Gumbel-Softmax mapping. Thus, our method addresses a different optimization problem and provides a mechanistic perspective on robustness through probabilistic reparameterization. Since acquisition protocol optimization and k-space subsampling are complementary, their integration is a promising direction for future work.

3. Methods

3.1. Physics-informed optimization of control parameter configurations

3.1.1. Overview for physics-informed optimization of control parameters

An algorithm for the physics-informed learning-based optimization of acquisition schedules (LOAS) is described in Appendix A. The MRI acquisition protocol consists of multiple configurations of control parameters. A deep neural network (θ) is designed to solve the inverse problem of physics-driven signal model for tissue characteristic estimation and thus the estimation errors with respect to the control parameters are calculated as follows (Kang et al., 2022; Perlman et al., 2022; Cohen & Otazo, 2023):

$$\mathcal{L}(\theta, \mathbf{cp}) = \left|\left|\widehat{tc} - tc\right|\right|_2^2 = \left|\left|f_\theta(\tilde{\mathbf{S}}(tc, \mathbf{cp})) - tc\right|\right|_2^2 \quad (1)$$

The backpropagation with gradient descent allows the control parameters to be updated to minimize estimation errors of tissue characteristics. For each iteration, the control parameters ($\mathbf{cp}$) and the inverse-problem-solving neural network (f_θ) are simultaneously updated.

3.1.2. Signal model governed by MRI physics

For tissue characteristic estimation, the signal evolutions obtained from various configurations of control (scan) parameters are fitted to the signal model. The physics-driven signal model is based on the principles of physics that explain the underlying mechanisms of the MRI system. Therefore, we can accurately simulate the MRI signal with given tissue characteristics and control parameters. Then, solving the inverse problem of the signal model ($\mathcal{M}$) can provide estimates of tissue characteristics. The MRI signal (S) is either determined by this physics model ($\mathcal{M}$) via simulation or acquired from an MRI system:

$$S(tc, cp) \equiv \mathcal{M}(tc, cp) \quad (2)$$

where tc is tissue characteristics and cp is a configuration of control parameters. Repeated, serial acquisitions are acquired with different configurations of control parameters to generate signal evolutions encoding tissue characteristics. Thus, a set of tissue characteristics results in a distinct signal evolution, effectively serving as its unique fingerprint. The signal evolution under given **cp** with respect to the tissue characteristics can be described as follows (Ma et al., 2013; Cohen-Adad et al., 2021; Jordan et al., 2021):

$$\begin{aligned} \mathbf{S}(tc, \mathbf{cp}) &= [S(tc, cp_1), \ldots, S(tc, cp_N)], \\ cp_i &= [cp_{i,1}, cp_{i,2}, \ldots, cp_{i,M}] \end{aligned} \quad (3)$$

where N represents the number of acquisitions and M is the number of control parameters for i^{th} acquisition. Although the inverse problem of physics-based model is often ill-posed due to its complexity, recent studies have addressed it with neural networks in a data-driven manner (Cohen et al., 2018; Aggarwal et al., 2018; Jun et al., 2021):

$$\widehat{tc} = \mathcal{M}^{-1}(\mathbf{S}(tc, \mathbf{cp})) = f_\theta(\mathbf{S}(tc, \mathbf{cp})) \quad (4)$$

where $\widehat{tc}$ is the estimated tissue characteristics by solving the inverse problem and f_θ is a deep neural network. However, the inverse problem becomes even more complicated in the real world due to noise (Chen et al., 2024; Kang et al., 2024; Ades-Aron et al., 2025). The noise could originate from various sources such as diverse physiological processes, thermal noise, motions, and many others (Brooks

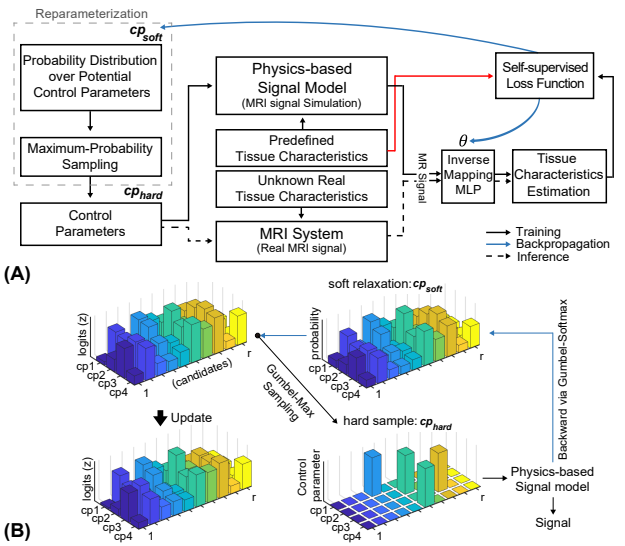

Figure 2. **(A)** Schematic comparison of the proposed RPO frameworks for control parameter optimization. **(B)** Each control parameter is represented by r candidates with associated probabilities. The candidate with the highest probability is selected via max function. These selected hard samples are then fed to a physics-based signal model to simulate the MR signal. By solving the inverse problem, the framework estimates the underlying tissue characteristics from the simulated signal. The resulting estimation error is backpropagated through the soft relaxation $\mathbf{cp}_{soft}$ using a straight-through Gumbel-Softmax estimator to update the learnable logits $\mathbf{z}$.

et al., 2013). Moreover, the control parameters themselves might introduce errors, i.e. the system may not execute the given control parameters exactly due to inherent systematic imperfection such as magnetic field inhomogeneity, eddy current, gradient nonlinearity, etc (Graves & Mitchell, 2013; Krupa & Bekiesińska-Figatowska, 2015; Manzano Patron et al., 2024; Kang et al., 2023; Simegn et al., 2025). These errors and noise should also be considered in the model as follows:

$$\widehat{tc} = \mathcal{M}^{-1}(\mathbf{S}(tc, \mathbf{cp}+\delta\mathbf{cp})+\mathbf{N}(\sigma)) = f_\theta(\tilde{\mathbf{S}}(tc, \mathbf{cp})) \quad (5)$$

where $\tilde{\mathbf{S}}$ is an observed signal evolution with system error and noise, $\delta\mathbf{cp}$ is the system error (e.g., software-hardware mismatch) for control parameters, and $\mathbf{N}$ is the noise originating from the system which is assumed to be additive white Gaussian noise with the noise level of σ in our study.

3.2. Robust Physics-informed Optimization (RPO) of control parameters

3.2.1. Probabilistic Reparameterization of Control Parameters

The overall architecture of proposed RPO method is described in Fig. 2. We use learnable unnormalized logits ($\mathbf{z}$) for each control parameter (cp) to have multiple candidate bins. Each bin corresponds to a respective candidate of control parameter and the number of bins determines the step size of control parameters. To obtain the maximum probability sample from the categorical distribution defined by the logits z, the Gumbel-Max trick is used to efficiently sample y using noise perturbation (Gumbel, 1954; Maddison et al., 2014). The control parameter with the highest probability is selected as the hard sample ($\mathbf{h}$):

$$\mathbf{h} = \text{one-hot}_r(\underset{l}{\text{argmax}}\,[z_l + g_l]) \quad (6)$$

where g_l is an i.i.d sample drawn from Gumbel (0, 1). The corresponding physical control parameter is obtained by projecting onto the predefined candidate vector as follows:

$$cp_{hard} = \mathbf{C}^\top \mathbf{h} = \sum_{l=1}^{r} C_l h_l \quad (7)$$

where cp_{hard} denotes the hard control parameter and $\mathbf{C} = \left[0, \frac{1}{r-1}, \frac{2}{r-1}, \ldots, 1\right]^T$ is a uniformly spaced candidate vector defined in a normalized $[0, 1]$ domain. The normalized control parameters are mapped to their physical ranges within the signal model. However, during optimization, the gradient cannot be computed when sampling from $\mathbf{z}$ because the max function is non-differentiable. We therefore use a straight-through (ST) Gumbel-Softmax estimator (Jang et al., 2017; Maddison et al., 2017), where $\mathbf{cp}_{hard}$ is used for the physics-based forward simulation and the continuous relaxation $\mathbf{cp}_{soft}$ serves as a differentiable surrogate for backpropagation (Fig. 2B):

$$\mathbf{cp}_{soft,l} = \mathbf{C}^\top \text{Softmax}_\tau(\mathbf{z}+\mathbf{g})_l = \mathbf{C}^\top \frac{e^{(z_l+g_l)/\tau}}{\sum_{n=1}^{r} e^{(z_n+g_n)/\tau}} \quad (8)$$

where $\mathbf{cp}_{soft}$ is the soft relaxation of the control parameter, and τ is the temperature of the Softmax operation, which determines the degree of relaxation from the discrete argmax operation to its continuous Softmax approximation. Thus, the unnormalized logits are updated via gradient descent:

$$\begin{aligned}\delta\mathbf{z} &= -\gamma_{cp}\frac{\partial\mathcal{L}}{\partial\mathbf{z}} \\ &= -\gamma_{cp}\frac{\partial\mathcal{L}}{\partial\mathbf{cp}_{\text{hard}}}\frac{\partial\mathbf{cp}_{\text{hard}}}{\partial\mathbf{z}} \approx -\gamma_{cp}\frac{\partial\mathcal{L}}{\partial\mathbf{cp}_{\text{hard}}}\frac{\partial\mathbf{cp}_{\text{soft}}}{\partial\mathbf{z}}\end{aligned} \quad (9)$$

where γ_{cp} denotes the learning rate for the control-parameter update. Thus, unlike LOAS, which directly updates control parameter values, RPO performs gradient-based updates on the logits of a categorical distribution, thereby optimizing control parameters through probabilistic reparameterization.

3.2.2. Theoretical analysis: noise robustness of gradient descent in RPO

To motivate the construction of the probabilistic representation of control parameters, we investigate the gradient of

the objective function ($\mathcal{L}$) with respect to $\mathbf{z}$ as follows:

$$\frac{\partial \mathcal{L}}{\partial \mathbf{z}} \approx \frac{\partial \mathcal{L}(\mathbf{S})}{\partial \mathbf{cp}_{hard}} \frac{\partial \mathbf{cp}_{soft}}{\partial \mathbf{z}} \tag{10}$$

In practice, we observe $\tilde{\mathbf{S}} = \mathbf{S} + \delta\mathbf{S}$, where $\delta\mathbf{S}$ captures noise and system imperfections. These perturbations induce discrepancies between the noiseless gradient and the estimated gradient, which can be interpreted as stochastic perturbations in stochastic gradient descent (SGD) dynamics (Li et al., 2020; Song et al., 2022; Li et al., 2017; Smith & Le, 2017). Let $\mathbf{x}$ denote the optimization variable generically, corresponding to the control parameters in LOAS and the logit variable $\mathbf{z}$ in RPO. Then, the update can be written as:

$$\mathbf{x}^{k+1} - \mathbf{x}^k = -\frac{\epsilon}{B}\nabla\mathcal{L}(\mathbf{x}^k) + \boldsymbol{\xi}, \tag{11}$$

where $\boldsymbol{\xi}$ represents the aggregated stochastic perturbation induced by measurement noise and finite-sample effects in SGD. Accordingly, we approximate the gradient with respect to $\mathbf{cp}_{hard}$ using a first-order perturbation expansion around the noiseless signal:

$$\frac{\partial \mathcal{L}(\tilde{\mathbf{S}})}{\partial \mathbf{cp}_{hard}} \approx \frac{\partial \mathcal{L}(\mathbf{S})}{\partial \mathbf{cp}_{hard}} + \mathbf{E}, \tag{12}$$

where $\mathbf{E}$ denotes the effective first-order gradient perturbation induced by signal noise, together with additional stochastic effects from measurement noise and finite-sample effects in SGD.

The proposed soft reparameterization induces the following decomposition of the gradient:

$$\begin{aligned}\frac{\partial \mathcal{L}(\tilde{\mathbf{S}})}{\partial \mathbf{z}} &\approx \frac{\partial \mathcal{L}(\tilde{\mathbf{S}})}{\partial \mathbf{cp}_{\text{hard}}} \frac{\partial \mathbf{cp}_{\text{soft}}}{\partial \mathbf{z}} \\ &= \left(\frac{\partial \mathcal{L}(\mathbf{S})}{\partial \mathbf{cp}_{\text{hard}}} + \mathbf{E}\right) \frac{\partial \mathbf{cp}_{\text{soft}}}{\partial \mathbf{z}} = \frac{\partial \mathcal{L}(\mathbf{S})}{\partial \mathbf{z}} + \mathbf{EJ}\end{aligned} \tag{13}$$

where $\mathbf{J} = \frac{\partial \mathbf{cp}_{soft}}{\partial \mathbf{z}}$ is the Jacobian of the Gumbel-Softmax mapping from $\mathbf{z}$ to $\mathbf{cp}_{soft}$. Notably, the softmax function is the gradient of the log-sum-exp function, whose Jacobian corresponds to its Hessian and is positive semidefinite (Gao & Pavel, 2017). Furthermore, the softmax mapping is globally contractive, with a Lipschitz constant of at most $1/2$ across ℓ_p norms (Nair, 2026). These properties imply that the Jacobian of the softmax mapping constrains the propagation of gradient perturbations.

Proposition 1 (Reparameterization-Induced Attenuation of Gradient Perturbations). *Let* $\mathbf{J}$ *denote the Jacobian of the Gumbel-Softmax mapping with temperature* τ*, and assume* $r \leq 43$*. Then, for any perturbation matrix* $\mathbf{E}$*,*

$$\|\mathbf{EJ}\|_2 < \frac{1}{\tau}\|\mathbf{E}\|_2 \tag{14}$$

This implies that the soft reparameterization attenuates gradient perturbations in operator norm. In particular, under our setting $\tau \geq 1$*, we have* $\|\mathbf{EJ}\|_2 < \|\mathbf{E}\|_2$ *(Proof in Appendix D).*

In contrast to conventional LOAS, where the gradient perturbation $\mathbf{E}$ directly affects the update in control space, RPO transforms the perturbation through the reparameterization Jacobian as $\mathbf{EJ}$. This provides a mechanistic explanation for the robustness of RPO: probabilistic relaxation does not merely enable differentiable optimization over discrete candidates, but also filters gradient perturbations induced by measurement and SGD noise.

4. Experiments

We simulated MR signals using a two-pool proton exchange physics model (*PM*) governed by modified Bloch-McConnell equations via pulseq-cest-library (Vladimirov et al., 2025) (See Appendix B). One million sets of predefined tissue characteristics, accounting for possible *in vivo* scenarios, were used for training. For the test data, ten thousand sets with a finer step size were used.

4.1. Monte Carlo simulation of gradient perturbation attenuation

4.1.1. Experiment setup

To validate the attenuation of gradient perturbations induced by the random interference term ($\mathbf{E}$) in Eq. 13, we compute the loss gradients with respect to each control parameter for both LOAS and RPO. We also evaluate multiple temperature values $\tau \in \{1, 2, 5, 10\}$ to validate the theoretical prediction in Proposition 1. We consider a configuration with four control parameters ($M = 4$) and ten acquisitions ($N = 10$). For a fair comparison, only the target control parameter is reparameterized via the soft relaxation, while the remaining three are kept unchanged. The gradient difference between the noisy (σ = 0.01) and noiseless systems was calculated using randomly initialized control parameters and inverse-problem-solving MLP, i.e., $mean|\partial\mathcal{L}(\tilde{\mathbf{S}})/\partial\mathbf{cp_i} - \partial\mathcal{L}(\mathbf{S})/\partial\mathbf{cp_i}|$, averaged over 10,000 random control parameter samples and test data. We tested two types of inverse-solving MLPs: a randomly initialized MLP for each run and a fixed MLP pre-trained with LOAS protocol. The mean absolute error (MAE) was computed by averaging the gradient differences over these samples.

4.1.2. Results

Table 1 summarizes the loss gradient differences between noisy and noiseless conditions. RPO consistently yields significantly smaller gradient discrepancies compared to LOAS, confirming that the perturbation term is effectively attenuated through the Jacobian of the Gumbel-Softmax

Table 1. The mean absolute error (MAE) of the loss gradient difference between noiseless and noisy conditions, computed across 10,000 random initializations of cp_i on test data.

	Random MLP (MAE $\times 10^7$)				Pre-trained MLP (MAE $\times 10^3$)			
	cp_1	cp_2	cp_3	cp_4	cp_1	cp_2	cp_3	cp_4
LOAS	2.88	5.17	1.49	0.42	11.29	13.58	3.50	1.18
RPO (τ=1)	0.53	0.64	0.22	0.11	0.77	0.44	0.22	0.15
RPO (τ=2)	0.22	0.28	0.09	0.04	0.40	0.22	0.11	0.06
RPO (τ=5)	0.06	0.09	0.03	0.01	0.15	0.08	0.04	0.02
RPO (τ=10)	0.03	0.04	0.01	0.01	0.08	0.04	0.02	0.01

mapping. Moreover, the results reveal a clear dependence on the temperature parameter τ: increasing τ progressively reduced gradient perturbations, consistent with the theoretical bound proposed in Proposition 1. These findings highlight the role of temperature in controlling the optimization dynamics: larger τ attenuates perturbations and stabilizes the updates while maintaining a broader distribution over candidates, whereas smaller τ sharpens the distribution and promotes a focused search during optimization. Notably, this behavior holds for both the pre-trained and randomly initialized MLPs, capturing the complete optimization trajectory from initialization to convergence.

4.2. Tissue characteristic estimation

4.2.1. Experiment setup

The goal of optimization of control parameters is to minimize the estimation error of tissue characteristics. Therefore, the estimation errors from RPO were compared to those from conventional methods under various conditions, including two noise types (Gaussian and Rician), noise levels (σ = 0, 0.01, and 0.02), and numbers of acquisitions (N = 10, 20, and 30).

We compared RPO to different scheduling schemes, namely, linear, random uniform (random), pseudo-random (PR) (Kim et al., 2020), correlation-based (Corr) optimization (Cohen & Rosen, 2017), Cramer-Rao lower bound (CRLB) based optimization (Zhao et al., 2018), and learning-based optimization of acquisition schedule (LOAS) (Kang et al., 2022). The implementation details for each scheduling scheme are described in Appendix C.

4.2.2. Training details

We used a multi-layer perceptron (MLP) with six hidden layers of 256 units each with ReLU activation for the estimation network. To enhance the speed and stability of optimization, the tissue characteristics were normalized to the [0, 1] range and Sigmoid activation was used for the last layer. We performed optimizations on a dataset comprising one million sets of four tissue characteristics, each of which was uniformly sampled from a predefined range (See Table 4). The Adam optimizer was employed to update the network and control parameters via back-propagation. Learning rates of inverse network (f_θ) and control parameters ($\mathbf{cp}$) were heuristically determined as follows: 10^{-4} and 10^{-4} for LOAS method and 10^{-4} and 10^{-2} for RPO method, respectively, within the range of 10^{-1} to 10^{-5}. The temperature for the soft-relaxed control parameters was annealed using the schedule τ = max (1.0, 10exp($-0.03t$)) of the iteration step t. The optimization was run for 100 iterations with a batch size of 256. The network was implemented using PyTorch on an NVIDIA RTX 3090.

4.2.3. Results

The estimation errors of tissue characteristics were evaluated across epochs under varying noise levels and acquisition settings (Fig. 6). As shown in Fig. 6, the inverse problem becomes more challenging, with fewer acquisitions or larger noise levels. The accuracy gap between the LOAS and RPO methods becomes larger as the complexity of inverse problem increases. In addition, LOAS exhibits high variance in estimation error, reflecting sensitivity to initialization and instability during optimization. In contrast, RPO consistently achieves lower variance, demonstrating more stable optimization through probabilistic search.

Table 2 further confirms that RPO outperforms all competing methods. Notably, even in the noiseless setting, RPO achieves lower errors than LOAS, suggesting that LOAS is more susceptible to stochastic gradient noise during optimization. Moreover, temperature annealing in RPO yields lower errors and variance compared to fixed-temperature settings, likely due to broader exploration in early iterations and reduced sensitivity to gradient perturbations (Fig. 3). Interestingly, CRLB achieves lower errors than LOAS, which contrasts with prior findings (Kang et al., 2022). This discrepancy is likely due to the smaller dataset used in our study (1M vs. 40M), highlighting the efficiency of RPO under data-limited conditions. Furthermore, RPO demonstrates improved computational efficiency, with optimization times of 2.3 hours for RPO, 2.7 hours for LOAS, and 84.3 hours for CRLB when the number of acquisitions (N) was set to 10 and the Rician noise level (σ) to 0.01.

In addition, discrete values of control parameters are more suitable for MRI systems, where hardware input values are

Table 2. NMSE for tissue characteristic estimation on a 10,000-sample test set. Each scheme was run with the acquisition number (N) of 10 and repeated eight times for reliability assessment. Values are reported as mean ± standard deviation.

Acquisition Protocols	NMSE (x100) for each noise level (σ)				
		Gaussian noise		Rician noise	
	$\sigma = 0$	$\sigma = 0.01$	$\sigma = 0.02$	$\sigma = 0.01$	$\sigma = 0.02$
Linear	1.83 ± 0.46	3.23 ± 0.24	3.85 ± 0.34	3.21 ± 0.23	3.84 ± 0.34
Random	0.20 ± 0.06	2.64 ± 0.18	3.47 ± 0.22	2.61 ± 0.19	3.44 ± 0.22
PR	0.12 ± 0.02	2.32 ± 0.37	3.16 ± 0.35	2.29 ± 0.38	3.13 ± 0.36
Corr	1.25 ± 0.69	3.36 ± 0.44	3.95 ± 0.39	3.36 ± 0.44	3.95 ± 0.38
CRLB	**0.08 ± 0.01**	1.65 ± 0.01	2.41 ± 0.02	1.61 ± 0.02	2.38 ± 0.02
LOAS	0.15 ± 0.04	1.61 ± 0.09	2.51 ± 0.16	1.62 ± 0.07	2.54 ± 0.20
RPO (τ=1)	0.13 ± 0.02	1.53 ± 0.03	2.15 ± 0.03	1.47 ± 0.02	2.12 ± 0.03
RPO	**0.08 ± 0.01**	**1.40 ± 0.01**	**2.10 ± 0.02**	**1.38 ± 0.02**	**2.06 ± 0.02**

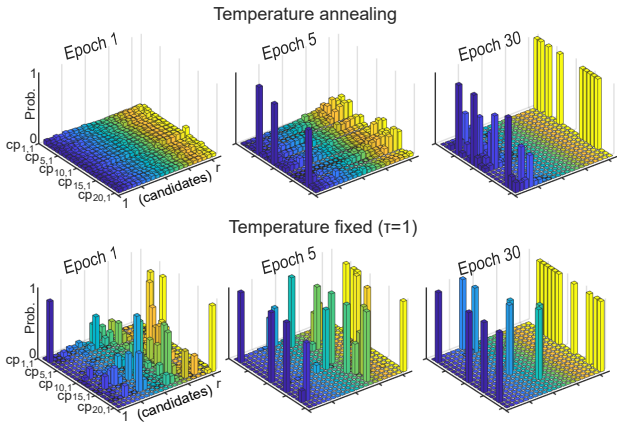

Figure 3. The probabilities for each candidate (π_l) of the first control parameter in i^{th} acquisition ($cp_{i,1}$) are displayed across different epochs. Here, $\Sigma_{l=1}^{r}\pi_l(cp_{i,1}) = 1$ and $cp_{i,1}$ are independent of each other. The acquisition number (N) of 20 and the noise level (σ) of 0.01 were used for RPO.

typically discretized, e.g., the minimum step size of radio frequency (RF) saturation time is often 50 ms due to its predefined block size of RF pulse and longer RF saturation is achieved by repetitively applying the RF blocks (Togao et al., 2016; Heo et al., 2019). This requires artificially rounding the control parameters from continuous to discrete space after optimization using conventional techniques, which could potentially undermine the optimization robustness (Wang et al., 2022; Ichikawa, 2024).

4.3. Probabilistic representation of control parameters

We assessed the probabilities associated with the soft-relaxed control parameters during the optimization process. As shown in Fig. 3, in the early iterations, the probabilities of all candidates are similar indicating that all possible control parameters are being equally considered. As iteration progresses, the probability of one candidate converges to one, while the probabilities of other candidates approach zero.

Table 3. MAE in the estimation of *in vivo* tissue characteristic maps.

	Tissue characteristics (tc)			
	tc_1 (Hz)	tc_2 $(\%)$	tc_3 (μs)	tc_4 (s)
RPO	**8.20**	**0.95**	**4.54**	**0.08**
LOAS	13.69	1.93	10.47	0.20
CRLB	14.69	1.36	7.35	0.19

On the other hand, RPO with fixed temperature converges faster to a certain value. Notably, the optimized values often lie at the lower or upper bounds of the predefined control parameter ranges, consistent with early observations (Zhao et al., 2018; Kang et al., 2022). This is commonly known as bang-bang behavior in optimal control problems and mitigating this may be an important direction for future studies (Seyde et al., 2021).

4.4. *In vivo* experiments

4.4.1. Experiment setup

All experiments were performed using a 3T MRI scanner (Achieva dStream, Philips Healthcare, Best, The Netherlands). A healthy volunteer dataset was acquired for a proof-of-concept *in vivo* feasibility study after written informed consent was obtained, in accordance with IRB requirements. The 3D ST images were scanned using a fat-suppressed, multi-shot turbo spin-echo (TSE) sequence with a two-channel parallel RF transmission technique. Detailed MRI settings are provided in Appendix E.2.1.

The acquisition protocol included a 60-dynamic reference scan, in which the increased number of independent dynamics leads to improved inverse problem accuracy and serves as a high-fidelity reference. In contrast, 10-dynamic scans were obtained using RPO-, LOAS-, and CRLB-optimized schedules. To assess repeatability, the optimization was re-run with different initializations, and the resulting schedules were used for subsequent data acquisition. For RPO and LOAS, the Rician noise level (σ) was set to 0.005.

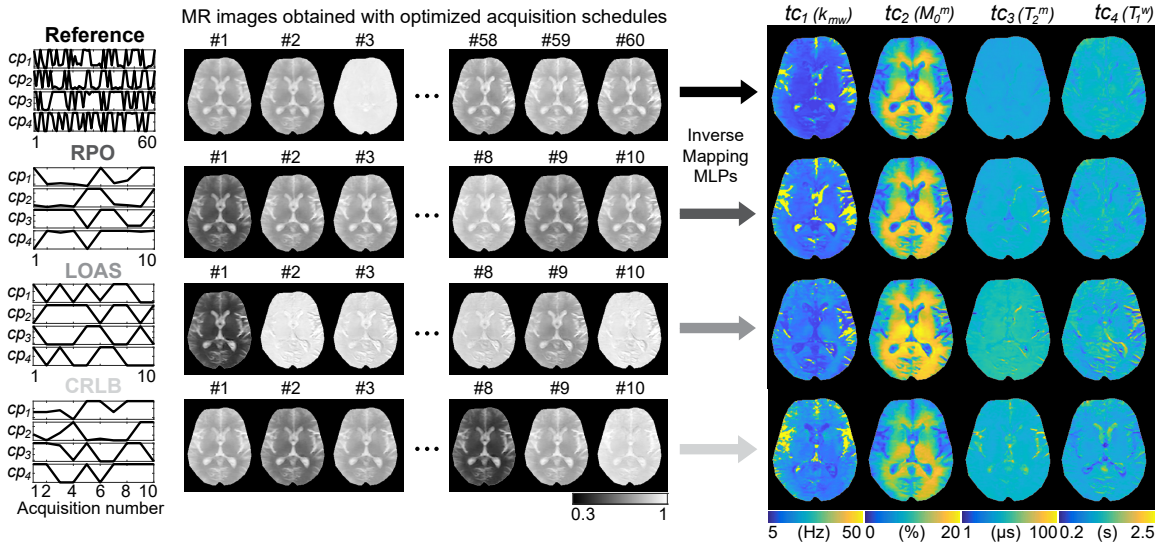

Figure 4. In vivo validation of 10-scan acquisition strategies (RPO, LOAS, and CRLB) against a 60-scan reference. Experimentally acquired MR signals are used as inputs to an MLP-based inverse mapping model, which generates quantitative maps of tissue characteristics.

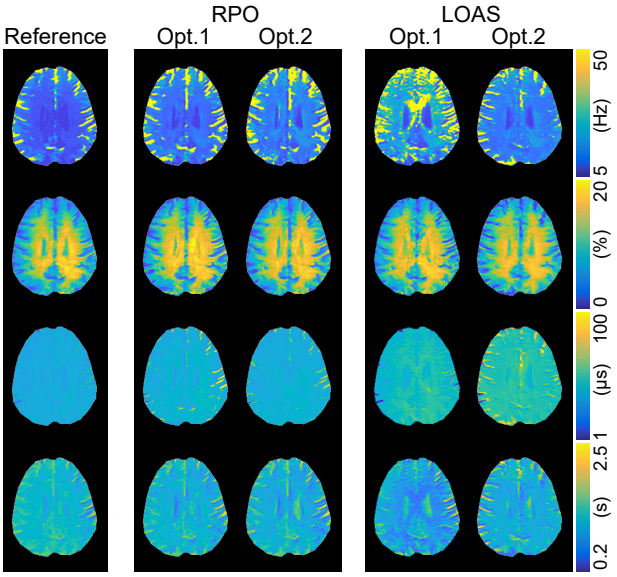

Figure 5. Evaluation of repeatability for RPO and LOAS optimization. The quantitative tissue maps generated by two independent optimization runs (Opt. 1 and Opt. 2) for each method.

Specifically, the system error of the control parameters was determined based on a real MRI system scenario, with $cp_{1,2}$ errors of 1% and $cp_{3,4}$ errors of 0.1%, respectively. The control-parameter-dependent noise conditions were also validated with a BrainWeb-based digital phantom (see Appendix E.3).

4.4.2. Results

The tissue characteristic maps reconstructed with the optimized acquisition protocol from RPO were in good agreement with the reference maps. In contrast, LOAS and CRLB showed poorer correspondence (Fig. 4). The averaged mean absolute errors (MAE) over 11 slices are shown in Table 3. RPO outperformed the LOAS and CRLB methods across all tissue characteristics. The *in vivo* data acquired from two different schedules optimized using the same strategy were used to estimate quantitative maps (Fig. 5). The RPO method showed higher consistency compared to the LOAS method, demonstrating high repeatability and stability of the optimization. Ultimately, RPO enables faster and more robust MR imaging, enhancing patient tolerance and mitigating motion artifacts. These improvements facilitate its transition into routine clinical practice.

5. Conclusion

We presented Robust Physics-informed Optimization (RPO), a probabilistic reparameterization framework for MRI protocol optimization. By reparameterizing control parameters with Gumbel-Softmax, RPO modifies how stochastic perturbations propagate during optimization, providing a mechanistic explanation for improved robustness to system and SGD noise. In addition, RPO maintains and updates a probabilistic distribution over the combinatorial control parameter space, enabling broader exploration than conventional point-estimate updates. Experiments on simulation, digital phantom, and *in vivo* MRI data show that RPO improves both tissue-estimation accuracy and optimization reliability over existing approaches. These results suggest that perturbation-aware reparameterization can be a useful design principle for physics-informed optimization of acquisition protocols and may extend to other signal models.

ACKNOWLEDGMENTS

This work was supported, in part, by the Institute of Information & Communications Technology Planning & Evaluation(IITP)-ITRC(Information Technology Research Center) grant funded by the Korea government(MSIT)(IITP-2026-RS-2024-00436857), by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT)(NRF-2026-RS-2024-00338025), and by grants from the National Institutes of Health (R01EB029974 and R01NS112242).

References

Ades-Aron, B., Coelho, S., Lemberskiy, G., Veraart, J., Baete, S. H., Shepherd, T. M., Novikov, D. S., and Fieremans, E. Denoising improves cross-scanner and cross-protocol test–retest reproducibility of diffusion tensor and kurtosis imaging. *Human Brain Mapping*, 46(4):e70142, 2025.

Aggarwal, H. K., Mani, M. P., and Jacob, M. MoDL: Model-based deep learning architecture for inverse problems. *IEEE transactions on medical imaging*, 38(2):394–405, 2018.

Brooks, J. C., Faull, O. K., Pattinson, K. T., and Jenkinson, M. Physiological noise in brainstem fMRI. *Frontiers in human neuroscience*, 7:623, 2013.

Chen, X., Wu, J., Yang, Y., Chen, H., Zhou, Y., Lin, L., Wei, Z., Xu, J., Chen, Z., and Chen, L. Boosting quantification accuracy of chemical exchange saturation transfer MRI with a spatial–spectral redundancy-based denoising method. *NMR in Biomedicine*, 37(1):e5027, 2024.

Cohen, O. and Otazo, R. Global deep learning optimization of chemical exchange saturation transfer magnetic resonance fingerprinting acquisition schedule. *NMR in biomedicine*, 36(10): e4954, 2023.

Cohen, O. and Rosen, M. S. Algorithm comparison for schedule optimization in MR fingerprinting. *Magnetic resonance imaging*, 41:15–21, 2017.

Cohen, O., Zhu, B., and Rosen, M. S. MR fingerprinting deep reconstruction network (DRONE). *Magnetic resonance in medicine*, 80(3):885–894, 2018.

Cohen, O., Young, R. J., and Otazo, R. Quantitative multislice and jointly optimized rapid CEST for in vivo whole-brain imaging. *Magnetic Resonance in Medicine*, 94(2):541–553, 2025.

Cohen-Adad, J., Alonso-Ortiz, E., Abramovic, M., Arneitz, C., Atcheson, N., Barlow, L., Barry, R. L., Barth, M., Battiston, M., Büchel, C., et al. Generic acquisition protocol for quantitative MRI of the spinal cord. *Nature protocols*, 16(10):4611–4632, 2021.

Collins, D. L., Zijdenbos, A. P., Kollokian, V., Sled, J. G., Kabani, N. J., Holmes, C. J., and Evans, A. C. Design and construction of a realistic digital brain phantom. *IEEE transactions on medical imaging*, 17(3):463–468, 1998.

Evans, M. R., Majumdar, S. N., and Schehr, G. Stochastic resetting and applications. *Journal of Physics A: Mathematical and Theoretical*, 53(19):193001, 2020.

Gao, B. and Pavel, L. On the properties of the softmax function with application in game theory and reinforcement learning. *arXiv preprint arXiv:1704.00805*, 2017.

Golub, G. H. and Van Loan, C. F. *Matrix computations*. JHU press, 2013.

Graves, M. J. and Mitchell, D. G. Body MRI artifacts in clinical practice: a physicist's and radiologist's perspective. *Journal of Magnetic Resonance Imaging*, 38(2):269–287, 2013.

Gumbel, E. J. *Statistical theory of extreme values and some practical applications: a series of lectures*, volume 33. US Government Printing Office, 1954.

Henkelman, R. M., Huang, X., Xiang, Q.-S., Stanisz, G. J., Swanson, S. D., and Bronskill, M. J. Quantitative interpretation of magnetization transfer. *Magnetic resonance in medicine*, 29(6): 759–766, 1993.

Heo, H.-Y., Zhang, Y., Lee, D.-H., Hong, X., and Zhou, J. Quantitative assessment of amide proton transfer (APT) and nuclear overhauser enhancement (NOE) imaging with extrapolated semi-solid magnetization transfer reference (EMR) signals: application to a rat glioma model at 4.7 tesla. *Magnetic resonance in medicine*, 75(1):137–149, 2016.

Heo, H.-Y., Zhang, Y., Lee, D.-H., Jiang, S., Zhao, X., and Zhou, J. Accelerating chemical exchange saturation transfer (cest) mri by combining compressed sensing and sensitivity encoding techniques. *Magnetic resonance in medicine*, 77(2):779–786, 2017.

Heo, H.-Y., Xu, X., Jiang, S., Zhao, Y., Keupp, J., Redmond, K. J., Laterra, J., van Zijl, P. C., and Zhou, J. Prospective acceleration of parallel RF transmission-based 3D chemical exchange saturation transfer imaging with compressed sensing. *Magnetic resonance in medicine*, 82(5):1812–1821, 2019.

Hernandez-Garcia, L., Aramendía-Vidaurreta, V., Bolar, D. S., Dai, W., Fernández-Seara, M. A., Guo, J., Madhuranthakam, A. J., Mutsaerts, H., Petr, J., Qin, Q., et al. Recent technical developments in asl: a review of the state of the art. *Magnetic resonance in medicine*, 88(5):2021–2042, 2022.

Hilbert, T., Xia, D., Block, K. T., Yu, Z., Lattanzi, R., Sodickson, D. K., Kober, T., and Cloos, M. A. Magnetization transfer in magnetic resonance fingerprinting. *Magnetic resonance in medicine*, 84(1):128–141, 2020.

Huang, T., Gong, Y.-J., and Zhang, J. Seeking multiple solutions of combinatorial optimization problems: A proof of principle study. In *2018 IEEE Symposium Series on Computational Intelligence (SSCI)*, pp. 1212–1218. IEEE, 2018.

Huang, T., Gong, Y.-J., Kwong, S., Wang, H., and Zhang, J. A niching memetic algorithm for multi-solution traveling salesman problem. *IEEE Transactions on Evolutionary Computation*, 24(3):508–522, 2019.

Ichikawa, Y. Controlling continuous relaxation for combinatorial optimization. *Advances in Neural Information Processing Systems*, 37:47189–47216, 2024.

Jang, E., Gu, S., and Poole, B. Categorical reparameterization with gumbel-softmax. In *International Conference on Learning Representations*, 2017.

Jian, S.-J. and Hsieh, S.-Y. A niching regression adaptive memetic algorithm for multimodal optimization of the euclidean traveling salesman problem. *IEEE Transactions on Evolutionary Computation*, 27(5):1413–1426, 2022.

Jordan, S. P., Hu, S., Rozada, I., McGivney, D. F., Boyacioğlu, R., Jacob, D. C., Huang, S., Beverland, M., Katzgraber, H. G., Troyer, M., et al. Automated design of pulse sequences for magnetic resonance fingerprinting using physics-inspired optimization. *Proceedings of the National Academy of Sciences*, 118(40):e2020516118, 2021.

Jun, Y., Shin, H., Eo, T., Kim, T., and Hwang, D. Deep model-based magnetic resonance parameter mapping network (DOPAMINE) for fast t1 mapping using variable flip angle method. *Medical Image Analysis*, 70:102017, 2021.

Jung, W., Bollmann, S., and Lee, J. Overview of quantitative susceptibility mapping using deep learning: Current status, challenges and opportunities. *NMR in Biomedicine*, 35(4):e4292, 2022.

Kang, B. and Seo, H. Prior-aware and context-guided group sampling for active probabilistic subsampling. In *The Fourteenth International Conference on Learning Representations*, 2026.

Kang, B., Kim, B., Schär, M., Park, H., and Heo, H.-Y. Unsupervised learning for magnetization transfer contrast MR fingerprinting: application to CEST and nuclear overhauser enhancement imaging. *Magnetic resonance in medicine*, 85(4): 2040–2054, 2021.

Kang, B., Kim, B., Park, H., and Heo, H.-Y. Learning-based optimization of acquisition schedule for magnetization transfer contrast MR fingerprinting. *NMR in Biomedicine*, 35(5):e4662, 2022.

Kang, B., Singh, M., Park, H., and Heo, H.-Y. Only-train-once MR fingerprinting for B0 and B1 inhomogeneity correction in quantitative magnetization-transfer contrast. *Magnetic Resonance in Medicine*, 90(1):90–102, 2023.

Kang, B., Lee, W., Seo, H., Heo, H.-Y., and Park, H. Self-supervised learning for denoising of multidimensional MRI data. *Magnetic resonance in medicine*, 92(5):1980–1994, 2024.

Kang, B., Singh, M., Seo, H., Park, H., and Heo, H.-Y. Physics-informed optimization of saturation-transfer MRI protocols using non-differentiable bloch models. *Physics in Medicine Biology*, 71(4):045001, 2026.

Kim, B., Schär, M., Park, H., and Heo, H.-Y. A deep learning approach for magnetization transfer contrast MR fingerprinting and chemical exchange saturation transfer imaging. *Neuroimage*, 221:117165, 2020.

Krupa, K. and Bekiesińska-Figatowska, M. Artifacts in magnetic resonance imaging. *Polish journal of radiology*, 80:93, 2015.

Kuppens, D., Barbieri, S., van den Berg, D., Schouten, P., Thoeny, H. C., van Laarhoven, H. W., Wennen, M., and Gurney-Champion, O. J. Acquisition-independent deep learning for quantitative MRI parameter estimation using neural controlled differential equations. *Medical Image Analysis*, 107:103768, 2026.

Lahiri, A., Fessler, J. A., and Hernandez-Garcia, L. Optimizing MRF-ASL scan design for precise quantification of brain hemodynamics using neural network regression. *Magnetic resonance in medicine*, 83(6):1979–1991, 2020.

Le Bihan, D. What can we see with IVIM MRI? *Neuroimage*, 187: 56–67, 2019.

Lee, W., Kim, B., and Park, H. Quantification of intravoxel incoherent motion with optimized b-values using deep neural network. *Magnetic Resonance in Medicine*, 86(1):230–244, 2021.

Li, M., Soltanolkotabi, M., and Oymak, S. Gradient descent with early stopping is provably robust to label noise for overparameterized neural networks. In *International conference on artificial intelligence and statistics*, pp. 4313–4324. PMLR, 2020.

Li, Q., Tai, C., et al. Stochastic modified equations and adaptive stochastic gradient algorithms. In *International Conference on Machine Learning*, pp. 2101–2110. PMLR, 2017.

Ma, D., Gulani, V., Seiberlich, N., Liu, K., Sunshine, J. L., Duerk, J. L., and Griswold, M. A. Magnetic resonance fingerprinting. *Nature*, 495(7440):187–192, 2013.

Maddison, C. J., Tarlow, D., and Minka, T. A* sampling. *Advances in neural information processing systems*, 27, 2014.

Maddison, C. J., Mnih, A., and Teh, Y. W. The concrete distribution: A continuous relaxation of discrete random variables. In *International Conference on Learning Representations*, 2017.

Manzano Patron, J. P., Moeller, S., Andersson, J. L., Ugurbil, K., Yacoub, E., and Sotiropoulos, S. N. Denoising diffusion MRI: Considerations and implications for analysis. *Imaging Neuroscience*, 2:1–29, 2024.

Meng, X., Cao, Y., and Zou, D. Per-example gradient regularization improves learning signals from noisy data. *Machine Learning*, 114(3):74, 2025.

Nair, P. Softmax is 1/2-lipschitz: A tight bound across all ℓ_p norms. *Transactions on Machine Learning Research*, 2026. ISSN 2835-8856.

Park, J., Jung, W., Choi, E.-J., Oh, S.-H., Jang, J., Shin, D., An, H., and Lee, J. DIFFnet: diffusion parameter mapping network generalized for input diffusion gradient schemes and b-value. *IEEE Transactions on Medical Imaging*, 41(2):491–499, 2021.

Perlman, O., Zhu, B., Zaiss, M., Rosen, M. S., and Farrar, C. T. An end-to-end ai-based framework for automated discovery of rapid CEST/MT MRI acquisition protocols and molecular parameter quantification (AutoCEST). *Magnetic Resonance in Medicine*, 87(6):2792–2810, 2022.

Perlman, O., Farrar, C. T., and Heo, H.-Y. MR fingerprinting for semisolid magnetization transfer and chemical exchange saturation transfer quantification. *NMR in Biomedicine*, 36(6): e4710, 2023.

Planchuelo-Gómez, Á., Descoteaux, M., Larochelle, H., Hutter, J., Jones, D. K., and Tax, C. M. Optimisation of quantitative brain diffusion-relaxation MRI acquisition protocols with physics-informed machine learning. *Medical Image Analysis*, 94:103134, 2024.

Scannell, C. M., Veta, M., Villa, A. D., Sammut, E. C., Lee, J., Breeuwer, M., and Chiribiri, A. Deep-learning-based preprocessing for quantitative myocardial perfusion MRI. *Journal of Magnetic Resonance Imaging*, 51(6):1689–1696, 2020.

Seyde, T., Gilitschenski, I., Schwarting, W., Stellato, B., Riedmiller, M., Wulfmeier, M., and Rus, D. Is bang-bang control all you need? solving continuous control with bernoulli policies. *Advances in Neural Information Processing Systems*, 34: 27209–27221, 2021.

Simegn, G. L., Sun, P. Z., Zhou, J., Kim, M., Reddy, R., Zu, Z., Zaiss, M., Yadav, N. N., Edden, R. A., van Zijl, P. C., et al. Motion and magnetic field inhomogeneity correction techniques for chemical exchange saturation transfer (CEST) MRI: A contemporary review. *NMR in Biomedicine*, 38(1): e5294, 2025.

Singh, M., Jiang, S., Li, Y., Van Zijl, P., Zhou, J., and Heo, H.-Y. Bloch simulator-driven deep recurrent neural network for magnetization transfer contrast MR fingerprinting and CEST imaging. *Magnetic resonance in medicine*, 90(4):1518–1536, 2023.

Singh, M., Kang, B., Mahmud, S. Z., van Zijl, P., Zhou, J., and Heo, H.-Y. Saturation transfer MR fingerprinting for magnetization transfer contrast and chemical exchange saturation transfer quantification. *Magnetic resonance in medicine*, 2025.

Smith, S. L. and Le, Q. V. A bayesian perspective on generalization and stochastic gradient descent. In *International Conference on Learning Representations*, 2017.

Smith, S. L., Dherin, B., Barrett, D. G., and De, S. On the origin of implicit regularization in stochastic gradient descent. In *International Conference on Learning Representations*, 2021.

Song, H., Kim, M., Park, D., Shin, Y., and Lee, J.-G. Learning from noisy labels with deep neural networks: A survey. *IEEE transactions on neural networks and learning systems*, 34(11): 8135–8153, 2022.

Togao, O., Hiwatashi, A., Keupp, J., Yamashita, K., Kikuchi, K., Yoshiura, T., Yoneyama, M., Kruiskamp, M. J., Sagiyama, K., Takahashi, M., et al. Amide proton transfer imaging of diffuse gliomas: effect of saturation pulse length in parallel transmission-based technique. *PloS one*, 11(5):e0155925, 2016.

Van Gorp, H., Huijben, I., Veeling, B. S., Pezzotti, N., and Van Sloun, R. J. Active deep probabilistic subsampling. In *International Conference on Machine Learning*, pp. 10509–10518. PMLR, 2021.

Van Zijl, P. C. and Yadav, N. N. Chemical exchange saturation transfer (cest): what is in a name and what isn't? *Magnetic resonance in medicine*, 65(4):927–948, 2011.

Vladimirov, N., Cohen, O., Heo, H.-Y., Zaiss, M., Farrar, C. T., and Perlman, O. Quantitative molecular imaging using deep magnetic resonance fingerprinting. *Nature Protocols*, pp. 1–31, 2025.

Wang, H. P., Wu, N., Yang, H., Hao, C., and Li, P. Unsupervised learning for combinatorial optimization with principled objective relaxation. *Advances in Neural Information Processing Systems*, 35:31444–31458, 2022.

Yoon, J., Gong, E., Chatnuntawech, I., Bilgic, B., Lee, J., Jung, W., Ko, J., Jung, H., Setsompop, K., Zaharchuk, G., et al. Quantitative susceptibility mapping using deep neural network: QSMnet. *Neuroimage*, 179:199–206, 2018.

Zhao, B., Haldar, J. P., Liao, C., Ma, D., Jiang, Y., Griswold, M. A., Setsompop, K., and Wald, L. L. Optimal experiment design for magnetic resonance fingerprinting: Cramér-rao bound meets spin dynamics. *IEEE transactions on medical imaging*, 38(3): 844–861, 2018.

A. Algorithm for Learning-based Optimization of Acquisition Schedule (LOAS)

All control parameters (cp) are optimized in a normalized $[0, 1]$ domain. A deterministic mapping within the signal model converts these normalized parameters to their physical ranges, thereby constraining the feasible parameter space based on prior knowledge of hardware constraints.

Algorithm 1 Physics-informed Optimization of Control Parameters

1. Randomly sample a batch of tissue characteristics (tc)
2. Simulate signal evolutions with sampled tissue characteristics (tc) and randomly initialized control parameters ($\mathbf{cp}$) via a physics-based signal model with random noise
3. Feed generated noisy signal evolutions ($\tilde{\mathbf{S}}$) into a neural network (f_θ) to estimate tissue characteristics
4. Compute the estimation loss:

$$\mathcal{L}(\theta, \mathbf{cp}) = \left\|\widehat{tc} - tc\right\|_2^2 = \left\|f_\theta(\tilde{\mathbf{S}}(tc, \mathbf{cp})) - tc\right\|_2^2$$

5. Update parameters via gradient descent:

$$\mathbf{cp}^{k+1} = \mathbf{cp}^k - \gamma_{cp}\frac{\partial \mathcal{L}}{\partial \mathbf{cp}}, \quad \theta^{k+1} = \theta^k - \gamma_\theta \frac{\partial \mathcal{L}}{\partial \theta}$$

6. Repeat until convergence

B. Bloch equations for two-pool proton exchange model

The behavior of spins in MRI system can be described with Bloch equations. Specifically, a two-pool proton exchange model, consisting of free bulk water pool (w) and the semisolid macromolecule pool (m), can be described with the modified Bloch-McConnell equations in the presence of radio frequency (RF) saturation. The signal originates from the protons in water so that the magnitude of net spin in longitudinal direction, which is equivalent to the longitudinal magnetization (M_z), in water pool is of interest (Henkelman et al., 1993; Heo et al., 2016):

$$M_z^w(t) = (M_0^w - M_{ss}^w)\, e^{\lambda t} + M_{ss}^w \tag{15}$$

$$M_{ss}^w = M_0^w \frac{\frac{1}{T_1^m}\left(k_{mw} M_0^m T_1^w\right) + \alpha}{\left(k_{mw} M_0^m T_1^w\right)\left(Absor_{RF,m} + \frac{1}{T_1^m}\right) + \alpha\left[1 + \left(\frac{\omega_1}{2\pi\Omega}\right)^2\left(\frac{T_1^w}{T_2^w}\right)\right]} \tag{16}$$

$$\lambda = -\frac{1}{2}\left(\alpha + \beta - \sqrt{(\alpha - \beta)^2 + 4k_{mw}^2 M_0^m}\right) \tag{17}$$

$$\alpha = \frac{1}{T_1^m} + k_{mw} + Absor_{RF,m} \tag{18}$$

$$\beta = \frac{1}{T_1^w} + k_{wm} + Absor_{RF,w} \tag{19}$$

$$Absor_{RF,i} = \frac{\omega_1^2 T_2^i}{1 + (2\pi\Omega T_2^i)^2} \tag{20}$$

where M_z^i is the longitudinal magnetization of a pool i; M_0^i is the equilibrium magnetization of a pool i; M_{ss}^w is the steady-state longitudinal magnetization of a pool i; T_1^i and T_2^i are the longitudinal and transverse relaxation times of a pool i, respectively; Ω is the frequency offset of the RF saturation; ω_1 is the RF saturation amplitude; k_{ij} is the proton exchange rate from a pool i to a pool j; and $Absor_{RF,i}$ is the RF absorption rate of a pool i. According to Eq.(15), the signal originating from the water pool is determined by three control parameters (cp): RF saturation power (B_1= $\omega_1/2\pi\gamma$; γ is the gyromagnetic ratio), frequency offset (Ω), and saturation time (t = T_{sat}). A relaxation delay time (Td) is additionally defined to consider the recovery of the longitudinal magnetization in the absence of RF saturation which determines the initial longitudinal magnetization for the next acquisition. Therefore, the final signal (S) can be described as follows:

$$S = \left[M_0^w \left(1 - e^{-Td/T_1^w}\right) - M_{ss}^w\right] e^{\lambda T_{sat}} + M_{ss}^w \tag{21}$$

$$tc = [k_{mw}, M_0^m, T_2^m, T_1^w], \quad cp = [B_1, \Omega, Ts, Td] \tag{22}$$

The ranges of tissue characteristics (tc) and control parameters (cp) are shown in Table 4. According to the step size of control parameters, the numbers of candidates are 21 for cp_1, 43 for cp_2, 33 for cp_3, and 21 for cp_4. The lower and upper bounds of control parameters were constrained by the hardware configurations, clinical limitations, and properties of tissues. For example, the limited range of the RF saturation power was used to stay within the clinically permitted specific absorption rate (SAR), mainly due to the use of the SAR-intensive, time-interleaved parallel transmission (pTX)-based RF saturation (Togao et al., 2016; Heo et al., 2019).

Table 4. Properties of control parameters (cp) and tissue characteristics (tc)

	cp_1 B_1 (μT)	cp_2 Ω (ppm)	cp_3 T_{sat} (s)	cp_4 T_d (s)	tc_1 k_{mw} (Hz)	tc_2 M_0^m (%)	tc_3 T_2^m (μs)	tc_4 T_1^w (s)
max	1.9	50	2.0	4.5	100	17	100	3.0
min	0.9	8	0.4	3.5	5	2	1	2.0
step size	0.05	1	0.05	0.05	1.0	0.1	1.0	0.03

Four tissue characteristics (tc) are exchange rate (k_{mw}), concentration (M_0^m), and transverse relaxation time (T_2^m) of macromolecule, and longitudinal relaxation time of free bulk water (T_1^w).

Four control parameters (cp) are RF saturation power (B_1), frequency offset (Ω), saturation time (T_{sat}), and relaxation delay time (T_d).

C. Implementation of comparison scheduling strategies

We performed several scheduling strategies for comparison. Conventional methods, which use indirect objective functions, do not simultaneously train the inverse-problem-solving MLP. Therefore, we separately trained an MLP for each schedule achieved from different strategies for each noise condition. Once optimal scheduling is achieved, a multi-layer perceptron (MLP) with six hidden layers of 256 units each with ReLU activation was trained for each schedule to enable estimation of tissue characteristics. We trained MLPs with one million sets of four tissue characteristics, each of which was uniformly sampled from a predefined range (Table 4). The networks were trained for 100 epochs with the adaptive moment estimation (ADAM) optimizer and a batch size of 256. The initial learning rate was 10^{-4} and decreased by a factor of 0.1 for every 30 epochs.

C.1. Linear

A linear schedule was generated by either linearly increasing values from the minimum to maximum points of each control parameter range or decreasing values from the maximum to minimum points (Table 4). We randomly selected whether to increase or decrease the values for each control parameter to form different linear schedules.

C.2. Random uniform (Random)

A random schedule was generated by randomly sampling control parameters from each control parameter range.

C.3. Pseudo-random (PR)

A PR schedule was generated by increasing temporal incoherence between acquisitions which is equivalent to reducing the information redundancy of repeated, serial acquisitions. For example, we selected four sets of two control parameters as follows: high cp_1 + high cp_2; high cp_1 + low cp_2; low cp_1 + high cp_2; low cp_1 + low cp_2. These combinations ensure the low information redundancy between acquisitions.

C.4. Correlation-based (Corr)

A correlation-based (Corr) schedule was generated by maximizing signal discrimination between different tissue types. The objective function was defined with the correlation coefficient between MR signals of the tissue characteristics (tc) dictionary, as follows (Cohen & Rosen, 2017):

$$min_{cp} \left|\left|I - D_{cp}^T D_{cp}\right|\right|_2^2 \tag{23}$$

where I is an identity matrix and D is a dictionary of MR signals consisting of 10,000 sets of tissue characteristics (tc). The optimization was solved using an interior-point (IP) algorithm.

C.5. Cramer-Rao lower bound (CRLB)

A CRLB-based schedule was generated by minimizing CRLB objective function of the physics model (PM). The CRLB function (V) was defined by the inverse of the Fisher information matrix under Gaussian noise assumption as follows:

$$Variance(tc) \geq V(tc) = I^{-1}(tc) \tag{24}$$

$$I(tc) = \frac{1}{\sigma^2}\Sigma_N(\frac{\partial \mathbf{S}(\mathbf{cp}, tc)}{\partial tc})^T(\frac{\partial \mathbf{S}(\mathbf{cp}, tc)}{\partial tc}) \tag{25}$$

where σ is the standard deviation of the noise, and N is the number of acquisitions. The optimization with CRLB function minimizes the variance of tc estimates, thus maximizing the SNR efficiency. Mathematically, such an optimal problem can be formulated as follows (Zhao et al., 2018):

$$min_{cp}\Sigma_1^k Tr(WV(tc)) \tag{26}$$

where k is the number of tissue characteristics sets used for optimization and W is a weighting function which was heuristically determined to normalize the effect of each tc. The above optimization was solved using a sequential quadratic programming (SQP) algorithm. The number of tc sets (k) was set to 10,000 for Corr and CRLB methods to align the computational time with that of RPO and LOAS methods.

D. Proof of Proposition 1.

Proof. Let $\mathbf{y} = \mathrm{Softmax}_\tau(\mathbf{z} + \mathbf{g})$. Then, the Jacobian of $\mathbf{y}$ with respect to $\mathbf{z}$ is given by the standard softmax Jacobian form (Gao & Pavel, 2017; Nair, 2026):

$$\mathbf{J}_y = \frac{\partial \mathbf{y}}{\partial \mathbf{z}} = \frac{1}{\tau}\left(\mathrm{diag}(\mathbf{y}) - \mathbf{y}\mathbf{y}^\top\right).$$

Define $\mathbf{M} = \mathrm{diag}(\mathbf{y}) - \mathbf{y}\mathbf{y}^\top$. Since the ℓ_∞-induced operator norm of a matrix is equal to its maximum absolute row sum (Golub & Van Loan, 2013), the ℓ_∞-induced operator norm of $\mathbf{M}$ is as follows:

$$\sum_{j=1}^{r} |M_{ij}| = y_i(1 - y_i) + \sum_{j \neq i} y_i y_j = 2y_i(1 - y_i).$$

Since $0 \leq y_i \leq 1$, the function $2y_i(1 - y_i)$ is maximized at $y_i = \frac{1}{2}$, we get

$$\|\mathbf{M}\|_\infty = \max_i \sum_{j=1}^{r} |M_{ij}| \leq \frac{1}{2}.$$

In addition, because $\mathbf{M}$ is symmetric, $\|\mathbf{M}\|_1 = \|\mathbf{M}\|_\infty$. Therefore, using the interpolation inequality $\|\mathbf{M}\|_2 \leq \sqrt{\|\mathbf{M}\|_1 \|\mathbf{M}\|_\infty}$ (Proposition 1 from (Nair, 2026)), we get the following bounds,

$$\|\mathbf{M}\|_2 \leq \frac{1}{2},$$

$$\|\mathbf{J}_y\|_2 = \left\|\frac{1}{\tau}\mathbf{M}\right\|_2 \leq \frac{1}{2\tau}. \tag{27}$$

Now, $\mathbf{cp}_{soft} = \mathbf{C}^\top \mathbf{y}$, so its Jacobian is

$$\mathbf{J} = \frac{\partial \mathbf{cp}_{soft}}{\partial \mathbf{z}} = \mathbf{C}^\top \mathbf{J}_y.$$

Since $\mathbf{J}_y \mathbf{1} = \mathbf{0}$ and $\mathbf{J}_y$ is symmetric, we also have $\mathbf{1}^\top \mathbf{J}_y = \mathbf{0}^\top$. Thus, for the centered vector $\tilde{\mathbf{C}} = \mathbf{C} - \bar{C}\mathbf{1}$, where $\bar{C} = \frac{1}{r}\sum_{l=1}^{r} C_l$, we obtain

$$\mathbf{J} = \mathbf{C}^\top \mathbf{J}_y = (\tilde{\mathbf{C}} + \bar{C}\mathbf{1})^\top \mathbf{J}_y = \tilde{\mathbf{C}}^\top \mathbf{J}_y.$$

Thus,

$$\|\mathbf{J}\|_2 \leq \|\tilde{\mathbf{C}}\|_2 \|\mathbf{J}_y\|_2.$$

For the uniform candidate vector $C_l = \frac{l-1}{r-1}$, we get

$$\|\tilde{\mathbf{C}}\|_2 = \sqrt{\frac{r(r+1)}{12(r-1)}} < 2 \quad (r \leq 43).$$

Therefore,

$$\|\mathbf{J}\|_2 < \frac{1}{\tau}. \tag{28}$$

Finally, for any perturbation matrix $\mathbf{E}$, we have

$$\|\mathbf{E}\mathbf{J}\|_2 \leq \|\mathbf{E}\|_2 \|\mathbf{J}\|_2 < \frac{1}{\tau}\|\mathbf{E}\|_2, \tag{29}$$

which shows that the mapping attenuates gradient perturbations.

E. Experiments

E.1. Tissue characteristic estimation with neural network

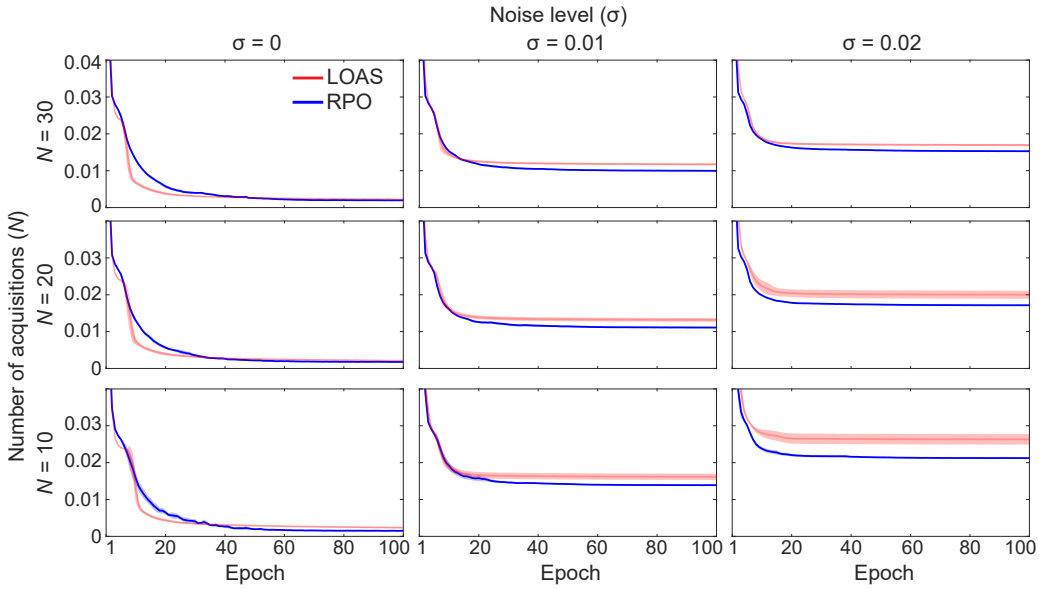

Figure 6. Train loss for tissue characteristic estimation as a function of epoch number, comparing LOAS and RPO approaches. Both optimizations were performed eight times to assess reliability, representing the 95% confidence interval with the shaded regions. Results are shown for various levels of Gaussian noise (σ) and number of acquisitions (N).

E.2. *in vivo* experiments

E.2.1. Implementation details: MRI acquisition settings

The imaging parameters were (Heo et al., 2017; Singh et al., 2023): TR = 6500 *ms*; TE = 6 *ms*; FOV = 212 × 192 × 36 mm^3; spatial resolution = 1.8 × 1.8 × 4 mm^3; turbo factor = 174; and slice sampling factor = 1.4. For compressed sensing, a variable density sampling pattern with a centric elliptical k-space ordering was used in the two phase-encoding directions (k_y-k_z) of the 3D acquisition, achieving a prospective acceleration factor of 4.

E.3. *In silico* tissue characteristic maps

E.3.1. Experiment setup

To evaluate the estimation on quantitative tissue maps, we make use of publicly available BrainWeb-based digital phantoms (Collins et al., 1998). Four tissue characteristic maps were generated with previously reported gray matter (GM) and white matter (WM) values (Kim et al., 2020). Synthetic MR images were generated from four tissue maps using optimized schedules and the physics model (Eq.(21)), then used to estimate tissue maps via an inverse-problem MLP. Here, the Gaussian noise level (σ) is set to 0.005. Specifically, the system error of the control parameters is determined based on a realistic MRI system scenario.

E.3.2. Results

The tissue characteristic maps reconstructed with the optimized schedule from RPO were in good agreement with the ground truth maps, whereas LOAS and CRLB showed unreliable results especially for $tc1$, which has an intrinsically low sensitivity (Figure 7). The averaged results over eight trials are shown in Table 5. The RPO outperforms LOAS and CRLB methods for all tissue characteristics. It is worth noting that any *in vivo* condition can be easily integrated into the forward physics model in RPO method, which provides high compatibility with various noise levels and types.

Table 5. Mean absolute errors (MAE) on estimation of *in silico* tissue characteristic maps. Each scheme was performed eight times with N of 20.

	Tissue characteristics (tc)			
	tc_1 (Hz)	tc_2 $(\%)$	tc_3 (μs)	tc_4 (s)
RPO	**0.90**	**0.50**	**0.91**	**0.026**
LOAS	1.35	0.57	1.07	0.027
CRLB	1.75	0.73	1.23	0.033

F. Limitations and Future Work

A potential limitation of this study is that the evaluation was conducted on a limited dataset acquired from a healthy volunteer. Further validation across diverse cohorts, patient populations, and scanner configurations is required to establish generalizability. Such evaluations will be pursued in future work.

Exploring alternative probabilistic search methods, such as Bayesian optimization (BO) or reinforcement learning (RL), could further strengthen our paper. However, to our knowledge, no previous studies have applied BO/RL to optimize acquisition schedules for quantitative MR. A fundamental challenge is that acquisition optimization must account for downstream inverse problem performance, which typically depends on a schedule-specific or pre-trained estimator. Therefore, BO or RL-based acquisition optimization typically requires repeated retraining of the inverse solver for each candidate schedule, unless an acquisition-independent estimator is available. Acquisition-agnostic quantification methods (Park et al., 2021; Kang et al., 2023; Kuppens et al., 2026) may serve as the objective function for BO/RL, making this an intriguing direction for future work.

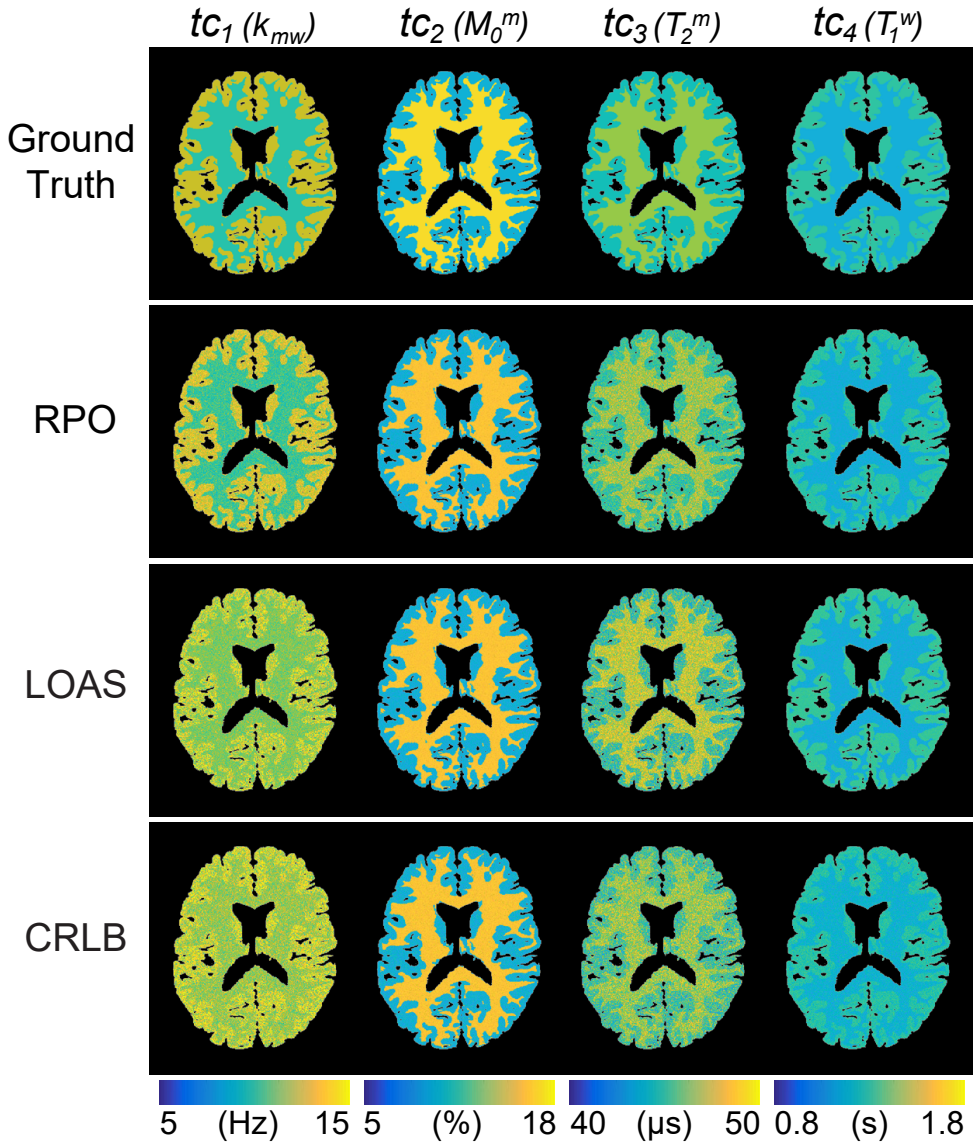

Figure 7. BrainWeb digital phantom studies with optimal acquisition scheduling from RPO, LOAS, and CRLB methods.