LEARNING CONTINUOUS ECHO TRAJECTORIES FOR SCAN-EFFICIENT QUANTITATIVE MRI

Anonymous authors
Paper under double-blind review

ABSTRACT

Quantitative MRI derives tissue parameters from how signals evolve across echo time, yet acquiring these signal samples increases scan burden. Can a retained source image support a continuously queryable echo trajectory with an interpretable decay structure? We address this question through cross-repetition-time echo-train completion, synthesizing a long-TR target train from one short-TR source echo per flip angle. The Endpoint-Conditioned Trajectory Network modulates source-image features, while the Physics-Constrained Decoder uses a shared reference amplitude and integrated positive rates to turn queried features into nonnegative, nonincreasing signals. The rate represents local log-signal decay, connecting continuous image generation to an explicit signal constraint. The model learns from echo images alone and supports TE queries beyond the supervised echo coordinates without retraining. On a private healthy-brain cohort, interior and terminal held-out-TE predictions remain close to their fully supervised counterparts. At comparable echo-image fidelity, the Physics-Constrained Decoder improves downstream agreement with acquired-echo-fitted T_2^* maps over the Unconstrained Decoder. The completion task targets omission of the long-TR block, corresponding to a nominal $4.67\times$ acceleration of the two-block repetition-time budget. These results connect continuous TE querying, interpretable decay, and acquisition-efficient synthesis within the evaluated protocol.

1 INTRODUCTION

Quantitative magnetic resonance imaging (qMRI) estimates tissue properties from signal changes across acquisition settings (Tofts, 2003; Deoni et al., 2003; Ma et al., 2013). For multi-echo imaging, the useful object is a signal trajectory over echo time (TE): its decay supports relaxation mapping, while acquiring the target echo train adds scan burden. We study cross-repetition-time (TR) completion in a MULTIPLEX (MTP) protocol (Ye et al., 2022), retaining one short-TR source echo per flip angle (FA) and synthesizing the long-TR target train. Omitting this target block would retain approximately 21.4% of the two-block repetition-time budget under equal repetition counts, motivating synthesis that supplies both target images and a coherent decay trajectory.

Missing-contrast networks can synthesize high-quality images (Dar et al., 2019; Dalmaz et al., 2022; Wu et al., 2026), and acquisition-conditioned models already support contrast queries (Xiong et al., 2024; 2025). Acquisition-conditioning alone, however, does not guarantee the signal ordering or decay structure needed for quantitative fitting. Under mono-exponential decay, T_2^* is determined by the negative reciprocal of the log-signal slope; small TE-dependent image errors can therefore alter the fitted map despite high per-echo fidelity. Our question is whether a continuously queryable representation with explicit decay structure can complete a missing acquisition block while improving downstream reference-map agreement.

We construct the predicted trajectory through two linked modules. The Endpoint-Conditioned Trajectory Network modulates subject-specific spatial features at each requested TE. The Physics-Constrained Decoder predicts a shared reference amplitude and a positive, TE-dependent rate, integrating the latter to generate the target train. This rate has the interpretation $r = -\partial_{\mathrm{TE}} \log S$: its positivity guarantees nonnegative, nonincreasing signals, while its TE dependence permits varying log-signal slopes. The Endpoint-Conditioned Trajectory Network supplies acquisition condition-

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

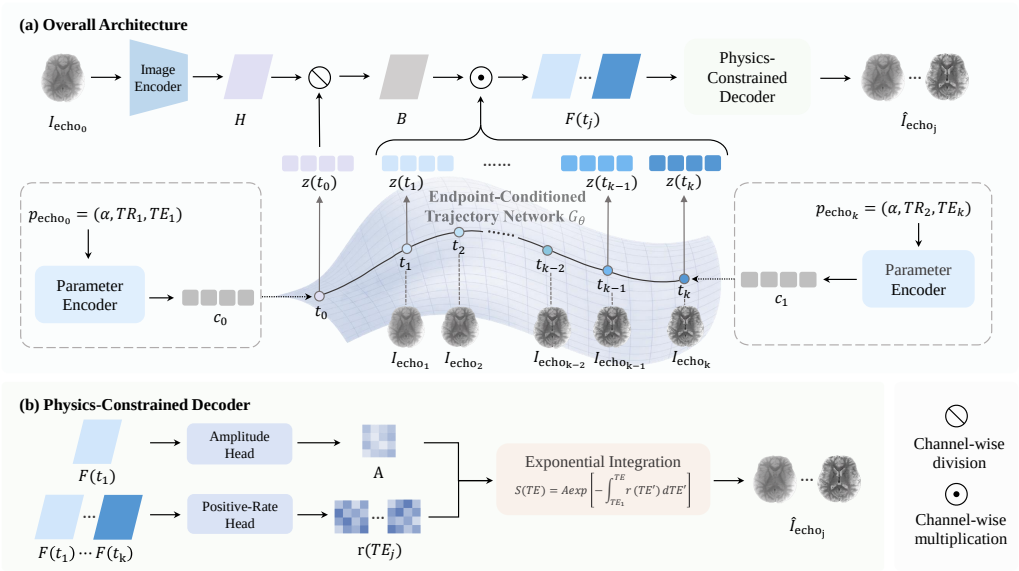


Figure 1: Framework overview. A retained source echo provides subject-specific spatial features, while numerical FA/TR/TE descriptors define endpoint-conditioned latent anchors. The Endpoint-Conditioned Trajectory Network produces continuously queryable trajectory states that modulate the factorized source representation. The Physics-Constrained Decoder predicts a reference amplitude and positive, TE-dependent decay rates, then generates the ordered target echo train through exponential integration.

ing; the Physics-Constrained Decoder supplies interpretable signal structure. Together they support user-specified TEs within the target interval without changing the trained output architecture.

A single source echo does not uniquely determine an individual's decay curve, so the missing information is inferred from the training population. The evaluation must therefore test whether the completed train retains subject-specific quantitative information, beyond producing plausible images or a decreasing signal. Image fidelity and downstream map agreement probe these complementary requirements. We evaluate echo reconstruction and agreement with maps fitted from acquired echoes, without quantitative-map supervision during training. Held-out experiments test prediction at acquired coordinates excluded from image supervision, while fully off-grid querying evaluates quantitative consistency under an unseen five-point TE grid.

Our contributions are:

- **Continuous echo generation with interpretable decay.** We combine the Endpoint-Conditioned Trajectory Network with the Physics-Constrained Decoder to generate a non-negative, nonincreasing target trajectory. A shared amplitude and local log-signal decay rates explicitly couple its queried echoes.
- **TE querying beyond supervised coordinates.** The trained model supports flexible sampling of the continuous target interval. Interior and terminal held-out-TE experiments evaluate prediction at coordinates excluded from image supervision, while fully off-grid five-echo querying evaluates quantitative consistency under an unseen TE grid.
- **Acquisition-efficient cross-TR completion.** One retained short-TR echo per FA supplies a synthesized long-TR train. The protocol implies a nominal 78.6% reduction of the two-block repetition-time budget under equal repetition counts, distinct from a measured scan-time reduction.
- **Downstream evidence beyond image fidelity.** A comparison with the Unconstrained Decoder shows improved acquired-reference T_2^* map agreement at comparable echo-image quality, supported by a subject-level paired analysis. This identifies a benefit of the Physics-Constrained Decoder within the evaluated protocol.

2 Related Work

Missing-contrast MRI synthesis. Deep learning has been widely studied for synthesizing missing MR contrasts from available acquisitions. Conditional adversarial models formulate the task as image-to-image translation between prescribed source and target contrasts (Dar et al., 2019), while residual vision transformers improve long-range representation learning for multimodal synthesis (Dalmaz et al., 2022). More recent dual-domain diffusion (Dayarathna et al., 2025) and discrete code-imputation frameworks (Wu et al., 2026) provide alternative mechanisms for improving image fidelity and handling missing observations. Structure-aware latent diffusion also supports synthesis under arbitrary missing-modality patterns within a unified model (Zhang et al., 2025). These methods recover a predefined target or target set. In the target-specific and one-to-many settings evaluated here, per-target fidelity does not explicitly constrain relations between adjacent echoes. We instead model the ordered echo train as a shared signal trajectory.

Continuous and monotonic neural representations. Implicit neural representations such as SIREN and NeRF encode a signal as a function of continuous coordinates (Sitzmann et al., 2020; Mildenhall et al., 2020). Neural ODEs model continuous hidden-state dynamics (Chen et al., 2018). Our Endpoint-Conditioned Trajectory Network similarly maps a continuous coordinate to a latent state. For ordered outputs, monotonic neural networks impose structure by construction: UMNN integrates a positive neural function (Wehenkel & Louppe, 2019), while Lipschitz monotonic networks combine weight constraints and a residual connection (Nolte et al., 2023). We use positive integration in the log-signal domain, conditioned on MRI acquisition descriptors and source-image features, to couple an entire target echo train. The resulting decay constraint is evaluated through downstream quantitative fitting.

Acquisition-conditioned synthesis. A complementary line of work represents image contrast through numerical acquisition parameters rather than discrete modality labels. CRL learns contrast representations conditioned on imaging parameters (Xiong et al., 2024), and CCR further disentangles contrast and anatomical content to support synthesis at flexible acquisition settings (Xiong et al., 2025). Related transformer frameworks likewise embed RF excitation parameters to generate multiple contrasts and quantitative maps (Nagar et al., 2026). We likewise use numerical acquisition descriptors, but generate the full TE-indexed echo train along a shared constrained trajectory.

Physics-informed learning for quantitative MRI. Embedding physical knowledge into learned models is a broader strategy for constraining forward and inverse problems (Karniadakis et al., 2021). Quantitative MRI methods often incorporate analytical signal models into reconstruction or parameter estimation. Joint MAPLE, for example, combines scan-specific learned reconstruction with joint quantitative parameter estimation (Heydari et al., 2024), illustrating the value of coupling image recovery with signal-model structure. Deep-learning superresolution has also been used to accelerate synthetic MRI with quantitative-map evaluation (Liu et al., 2026). Our setting completes a missing echo train and fits T_2^* only downstream, without map supervision. Acquisition physics constrains synthesis through positive decay rates integrated over TE, preserving cross-echo structure for subsequent fitting. Direct synthesis from fingerprinting or multitasking representations (Wang et al., 2023; Qiu et al., 2023) assumes a dedicated fingerprinting or multitasking acquisition as input, whereas our setting completes a partially acquired echo train from a retained source echo.

3 Method

We represent the missing target train as a continuous signal function conditioned on the retained source image and acquisition descriptors. The construction separates three roles: source features retain spatial information, an endpoint-conditioned latent function specifies TE-dependent modulation, and positive-rate integration couples the signals through an interpretable decay process. The structural constraint is imposed on the generated signal, rather than inferred from the latent path alone. An acquisition is described by $\mathbf{p} = (\alpha, \mathrm{TR}, \mathrm{TE}) \in \mathcal{P}$, where α denotes the flip angle (FA), TR the repetition time, TE the echo time, and $\mathcal{P}$ is the acquisition-parameter domain of interest. For each acquisition group $g \in \mathcal{G}$, a retained source image $I^{(g)}_{\mathrm{echo}_0}$ and its descriptors condition a set of target observations $\{(\mathrm{TE}^{(g)}_j, I^{(g)}_{\mathrm{echo}_j})\}_{j=1}^{K_g}$. Here TE denotes physical echo time, whereas t de-

notes a dimensionless path coordinate. The objective is to learn a representation from which echoes can be queried at a set of M_g locations, rather than restricting the output count to the K_g training observations.

Figure 1 summarizes the source-feature modulation and decay-constrained decoding pipeline. We use $t \in [t_0, t_K]$ as the nominal global path domain, with the current target train occupying $[t_1, t_K]$. Within a fixed target acquisition block, TE queries are mapped by

$$t(\mathrm{TE}) = t_1 + (t_K - t_1)\frac{\mathrm{TE} - \mathrm{TE}_1}{\mathrm{TE}_K - \mathrm{TE}_1}. \tag{1}$$

The source state is indexed by t_0, while t_1 and t_K denote the first and terminal target states. We normalize $t_0 = 0$ and $t_K = 1$; protocol-specific TE values and query coordinates are given in Section 4. The source condition defines the starting acquisition anchor at $t = t_0$, whereas signal integration starts at the first *target* echo at $t = t_1$, not at the source acquisition. We suppress group superscripts below except in the training objectives.

3.1 Source Features and Acquisition Anchors

An image encoder extracts a spatial feature field from the normalized source image:

$$\mathbf{H} = E_{\mathrm{img}}(I_{\mathrm{echo}_0}), \qquad \mathbf{H}(x) \in \mathbb{R}^C, \tag{2}$$

where $x \in \Omega$ indexes a spatial location and C is the feature dimension. A shared parameter encoder maps the source and terminal target acquisition descriptors to latent anchors:

$$\mathbf{c}_0 = E_{\mathrm{param}}(\mathbf{p}_{\mathrm{echo}_0}), \qquad \mathbf{c}_1 = E_{\mathrm{param}}(\mathbf{p}_{\mathrm{echo}_K}), \qquad \mathbf{c}_0, \mathbf{c}_1 \in \mathbb{R}^C. \tag{3}$$

The terminal anchor is determined by the known acquisition descriptors alone, not an additional target image.

Let $\mathbf{z}(t_0)$ denote the starting state of the latent path defined in Section 3.2. We factorize the source features relative to this state:

$$\mathbf{B}(x) = \frac{\mathbf{H}(x)}{d_\epsilon\big(\mathbf{z}(t_0)\big)}, \tag{4}$$

where $\mathbf{H}$ is the encoder output in Eq. (2), division is element-wise, the latent vector is broadcast spatially, and d_ϵ is defined in Appendix B.3.

The resulting $\mathbf{B}$ is a learned, subject-specific spatial representation held fixed over the target train.

3.2 Continuous Endpoint-Conditioned Feature Modulation

The Endpoint-Conditioned Trajectory Network models acquisition-dependent variation through a shared, endpoint-conditioned latent path:

$$\mathbf{z}(t) = G_\theta\big([\mathbf{c}_0, \mathbf{c}_1, t]\big), \qquad \mathbf{z}(t) \in \mathbb{R}^C, \tag{5}$$

where $[\cdot]$ denotes concatenation. For fixed acquisition anchors, varying t defines a continuous latent path shared by the requested echoes. Because the anchors encode descriptors alone, subjects acquired with the same protocol share this path; subject dependence enters through $\mathbf{B}(x)$. We implement G_θ as a multilayer perceptron with ReLU hidden activations and a linear output. Here "trajectory" denotes this coordinate-indexed latent function. Endpoint descriptors are fixed within each FA group in our experiments, so the conditioning comparison evaluates alternative parameterizations within that protocol, not generalization to unseen acquisition endpoints.

For a query at t, the path state modulates the factorized spatial features:

$$\mathbf{F}(x, t) = \mathbf{B}(x) \odot \mathbf{z}(t) = E_{\mathrm{img}}(I_{\mathrm{echo}_0})(x) \odot \frac{\mathbf{z}(t)}{d_\epsilon(\mathbf{z}(t_0))}, \tag{6}$$

where $\odot$ and the ratio act channel-wise.

3.3 Interpretable Decay through Positive-Rate Integration

The Physics-Constrained Decoder uses two voxelwise heads to map the trajectory-conditioned features to a nonnegative reference amplitude at the first target echo and a positive, TE-dependent decay rate:

$$A(x) = D_A(\mathbf{F}(x, t_1)), \tag{7}$$

$$r(x, \mathrm{TE}) = \mathrm{softplus}\left(D_r(\mathbf{F}(x, t(\mathrm{TE})))\right). \tag{8}$$

Here D_A denotes the amplitude head including its nonnegative output mapping, and the softplus rate parameterization ensures $r(x, \mathrm{TE}) > 0$. The amplitude is referenced to TE_1 and shared across all queries within the target train, and $r(x, \mathrm{TE})$ is expressed in ms^{-1}.

The generated signal is parameterized by exponential integration:

$$\widehat{S}(x, \mathrm{TE}) = A(x) \exp\left[-\int_{\mathrm{TE}_1}^{\mathrm{TE}} r(x, \mathrm{TE}')\, \mathrm{dTE}'\right]. \tag{9}$$

The synthesized echo image at query TE_j is $\widehat{I}_{\mathrm{echo}_j}(x) = \widehat{S}(x, \mathrm{TE}_j)$. A constant r recovers mono-exponential decay, while the predicted TE-dependent rate permits a more flexible nonincreasing signal family. For $A(x) > 0$, the rate satisfies $r(x, \mathrm{TE}) = -\partial_{\mathrm{TE}} \log \widehat{S}(x, \mathrm{TE})$, so the constraint has a direct interpretation as a nonnegative local log-signal decay. It models the expected decay envelope; noisy acquired magnitudes may fluctuate around that envelope (Gudbjartsson & Patz, 1995).

The rate interpretation applies to the synthesized signal; it does not identify a uniquely recovered tissue parameter. The constraint acts within each fixed-FA, fixed-TR target train. Cross-TR amplitude changes are learned from paired images, and a common relaxation parameter across FAs is not explicitly enforced.

For ordered TE queries starting at TE_1, we evaluate the integral using cumulative trapezoidal integration. Let $\widehat{S}_j(x)$ denote the resulting signal at the jth query. Because the amplitude head gives $A(x) \geq 0$ and the rate parameterization gives $r(x, \mathrm{TE}) > 0$, accumulated decay is nondecreasing for ordered TEs, yielding

$$0 \leq \widehat{S}_{j+1}(x) \leq \widehat{S}_j(x). \tag{10}$$

The implementation integrates over the supplied query grid, always starting at TE_1. Thus the discrete output also depends on the quadrature grid; Appendix B.2 verifies that this numerical dependence is small for uniform-dense and nonuniform alternatives at the common TE coordinates.

Flexible TE sampling. Equations (5)–(9) accept any finite ordered set of TE queries in the target interval, with the reference TE included for integration.

3.4 Learning from Echo Images

Let $\mathcal{K}_g \subseteq \{1, \ldots, K_g\}$ be the image-supervised targets in group g, and let $N_{\mathrm{sup}} = \sum_g |\mathcal{K}_g|$. We train on normalized echo images with

$$\mathcal{L}_{\mathrm{E}} = \frac{1}{N_{\mathrm{sup}}} \sum_{g \in \mathcal{G}} \sum_{j \in \mathcal{K}_g} \left\| \widehat{I}^{(g)}_{\mathrm{echo}_j} - I^{(g)}_{\mathrm{echo}_j} \right\|_{1,\mathrm{mean}}, \tag{11}$$

where $\|\cdot\|_{1,\mathrm{mean}}$ denotes the mean absolute entry value. The loss is averaged over training examples in each minibatch. Holding out a TE removes its target image from $\mathcal{K}_g$, without requiring a change to the network's query interface.

Boundary alignment connects the global path endpoint states to the acquisition anchors:

$$\mathcal{L}_{\mathrm{B}} = \|\mathbf{z}(t_0) - \mathbf{c}_0\|_1 + \|\mathbf{z}(t_K) - \mathbf{c}_1\|_1. \tag{12}$$

This soft latent-space constraint is evaluated for each acquisition group and averaged over the minibatch. The final training objective is

$$\mathcal{L} = \mathcal{L}_{\mathrm{E}} + \lambda_B \mathcal{L}_{\mathrm{B}}, \qquad \lambda_B = 0.1. \tag{13}$$

All synthesis components are optimized jointly. Training is confined to the observed target interval ($t \leq t_K$) and does not use T_2^* map supervision.

The completed echo train is subsequently passed to a separate T_2^* fitting procedure for quantitative evaluation. That fitting procedure is not part of the synthesis training graph.

4 EXPERIMENTS

We evaluate three aspects of the proposed representation: completion of the omitted long-TR block, prediction at TEs excluded from image supervision, and the effect of decoder structure on downstream fitting. Dense-query analyses additionally examine continuous signal behavior and numerical stability.

4.1 ACQUISITION-REDUCTION TASK AND EVALUATION SETUP

Data and acquisition. We use a private MTP brain MRI dataset from 42 healthy participants, with 32 for training and ten for evaluation. The MULTIPLEX (MTP) gradient-echo protocol (Ye et al., 2022) varies the excitation flip angle (FA), the interval between excitations (TR), and the time at which an echo is sampled after excitation (TE). Each FA (4° or 16°) has two acquisition blocks: $\mathrm{TR}_1 = 7.97$ ms and $\mathrm{TR}_2 = 29.23$ ms.

Completion task. We retain one TR_1 source echo at $\mathrm{TE} = 3.86$ ms per FA. A shared model predicts five TR_2 target echoes at TE $\{3.86, 8.95, 14.04, 19.13, 24.22\}$ ms, giving two source images and ten targets per subject. Each target train varies TE at fixed FA and TR. Source and target amplitudes can differ because they belong to different TR blocks. Both completed trains are jointly fitted to produce one T_2^* map, with the acquired trains supplying its reference. The omitted TR_2 block is a multi-echo acquisition in which all five target echoes are read within a single TR_2 repetition.

Acquisition budget. Under equal repetition counts, retaining the short-TR block gives a repetition-time budget of $7.97/(7.97 + 29.23) = 0.214$, a nominal 78.6% reduction or $4.67\times$ acceleration of that component. This protocol-based estimate excludes scan overheads and is not a measured wall-clock speedup.

For this protocol, the five target path coordinates are $\{0.2, 0.4, 0.6, 0.8, 1\}$.

Training and evaluation. All models use a fixed 1000-epoch training budget and the final-epoch checkpoint. The evaluation subjects are used only for final assessment, not training or model selection. We compare with ResViT (Dalmaz et al., 2022), D2Diff (Dayarathna et al., 2025), CodeBrain (Wu et al., 2026), and CCR (Xiong et al., 2025). Image fidelity uses PSNR and SSIM (Wang et al., 2004); quantitative agreement uses map MAE/R^2: the acquired and all synthesized echo trains are fitted with one common implementation, and every method is scored on the same fixed acquired-reference mask. Metrics are summarized across ten subjects as mean±SD.

4.2 FROM ECHO-TRAIN COMPLETION TO QUANTITATIVE MAPPING

Table 1: Image-level echo synthesis and downstream T_2^* map agreement over ten subjects (mean±SD). PSNR/SSIM average five target echoes per FA; map metrics evaluate the joint fit across both FAs.

	Mean echo-image quality				Echo-fitted T_2^* map quality	
	FA 4°		FA 16°		Joint fit across both FAs	
Method	PSNR↑	SSIM↑	PSNR↑	SSIM↑	MAE↓	R^2 ↑
ResViT	32.69 ± 1.11	0.9370 ± 0.0084	32.51 ± 1.31	0.9422 ± 0.0085	15.62 ± 2.27	0.191 ± 0.041
D2Diff	34.97 ± 1.02	0.9510 ± 0.0067	34.89 ± 1.15	0.9535 ± 0.0067	15.56 ± 1.84	0.196 ± 0.045
CodeBrain	31.88 ± 0.97	0.9286 ± 0.0086	31.95 ± 1.09	0.9351 ± 0.0082	16.12 ± 2.03	0.191 ± 0.033
CCR	34.66 ± 0.96	0.9471 ± 0.0069	34.66 ± 1.16	0.9511 ± 0.0070	15.40 ± 2.09	0.178 ± 0.044
Ours	$\mathbf{35.21} \pm 0.96$	$\mathbf{0.9514} \pm 0.0066$	$\mathbf{35.52} \pm 1.15$	$\mathbf{0.9577} \pm 0.0064$	$\mathbf{14.13} \pm 1.61$	$\mathbf{0.318} \pm 0.088$

Image metrics are averaged across echoes within subjects, then across subjects. Appendix Tables 3–4 report all five echoes separately. Our model achieves the highest mean PSNR in both FA groups

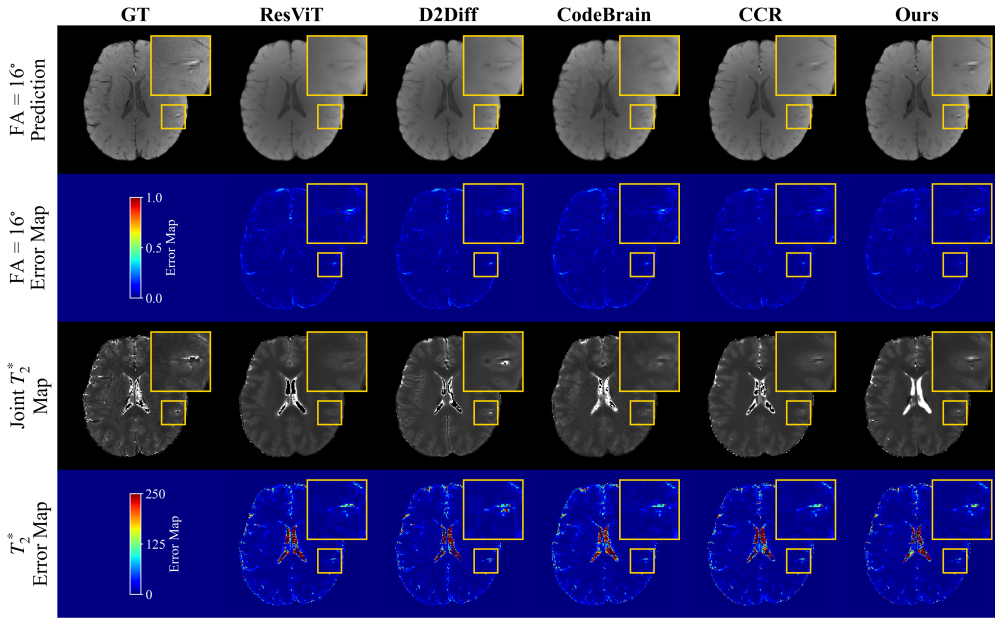


Figure 2: Qualitative echo completion and fitted-map comparison on subject 06, slice 6. The first two rows show FA $= 16°$ synthesis at the longest acquired echo, TE $= 24.22$ ms, and its absolute error. The last two rows show the jointly fitted T_2^* maps and corresponding errors. The acquired-reference column shows the ground-truth target echo to which every prediction and every error map is compared. Echo error is displayed over $[0, 1]$ and T_2^* error over $[0, 250]$ ms.

(35.21 and 35.52 dB), exceeding D2Diff by 0.24 and 0.63 dB, respectively. SSIM also ranks highest at $0.9514/0.9577$.

Figure 2 shows subject 06 at the longest acquired TE, where decreased signal and greater contrast difference from the source make reconstruction challenging. Under shared display ranges, our prediction better retains local structure and suppresses large residuals; the lower rows show the corresponding echo-fitted T_2^* maps.

The T_2^* columns of Table 1 evaluate whether image-level accuracy is accompanied by preservation of the cross-TE relationships used for quantitative estimation. Under the fixed acquired-reference mask and common scoring protocol, our completed trains yield the lowest MAE of 14.13 ± 1.61 ms and the highest map-agreement R^2 of 0.318 ± 0.088 among the evaluated methods. Supplementary tissue- and atlas-region analyses are reported in Appendix D.

4.3 TE Querying Beyond Supervised Coordinates

Held-out acquired coordinates. We withhold either the interior TE $= 14.04$ ms or the terminal TE $= 24.22$ ms, excluding the corresponding images from all training objectives and checkpoint selection while retaining their numerical acquisition descriptors. The former tests interpolation within the image-supervised range; the latter queries the acquisition boundary beyond the last supervised TE at 19.13 ms. Figure 3 shows that held-out PSNR remains within 0.05 dB of full supervision for both FAs. Under matched hold-outs, the Full Model exceeds Pointwise MLP by 0.38–0.52 dB and yields better fitted-map MAE/R^2 (Appendix Table 7).

At the interior TE, held-out PSNR reaches $34.01/34.04$ dB for FA $4°/16°$, respectively. Its jointly fitted map achieves MAE 14.17 ± 1.58 ms and $R^2 = 0.306 \pm 0.083$, compared with 14.13 ± 1.61 ms and 0.318 ± 0.088 under full supervision. The terminal hold-out yields MAE 14.28 ± 1.59 ms and $R^2 = 0.295 \pm 0.080$. At the terminal TE, held-out predictions reach 30.52 ± 1.26 and 31.84 ± 1.43 dB for the two FAs, respectively; SSIM differs from full supervision by at most 0.0002. Neither test evaluates unseen FA, TR, or endpoint descriptors.

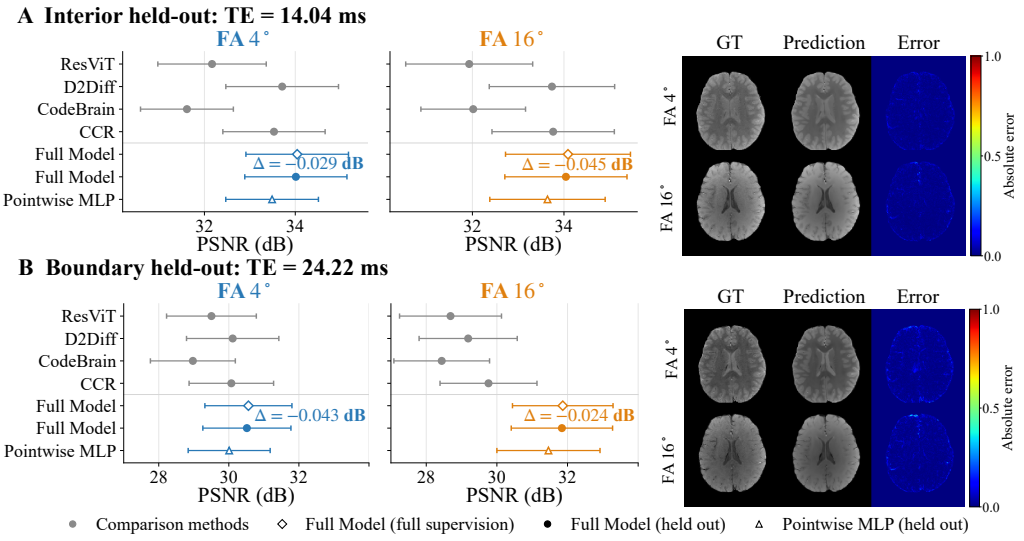


Figure 3: Prediction at two acquired TE coordinates withheld from image supervision. (A) Interior hold-out at 14.04 ms and (B) boundary hold-out at 24.22 ms. Plots report FA-specific PSNR over ten subjects (mean±SD); montages show the acquired target, held-out prediction, and absolute error for subject 06, slice 6. Gray comparison methods receive image supervision at all five target echoes and are shown only as fully supervised image-quality references; they are not trained under the held-out protocol. The annotated Δ is held-out minus full-supervision PSNR.

Fully off-grid TE querying. We further query five uniformly spaced echo times at TE $= \{5, 9, 13, 17, 21\}$ ms. All lie within the target TE interval, and none coincides with a supervised TE. Figure 4A shows the resulting off-grid echo trains for both FAs, and Figure 4B compares the acquired-reference T_2^* map with maps obtained from the standard-grid and off-grid synthesized echo trains.

The standard-grid synthesized echoes yield a T_2^* MAE of 14.13 ± 1.61 ms and $R^2 = 0.318 \pm 0.088$. Querying the entirely off-grid five-point set gives a closely comparable MAE of 14.24 ± 1.62 ms and $R^2 = 0.305 \pm 0.096$. These results show that the learned trajectory can be resampled on an unseen in-range TE grid while largely preserving downstream quantitative agreement. Unlike the acquired-coordinate hold-outs, this test evaluates quantitative consistency after replacing all five target sampling times.

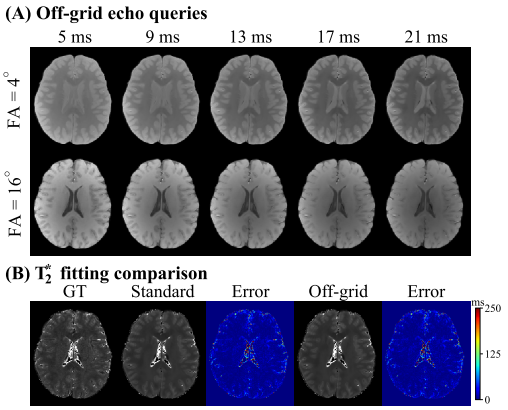


Figure 4: Off-grid TE querying and quantitative consistency for subject 06, slice 6. (A) Off-grid echoes at five unseen TEs for both FAs. (B) Joint T_2^* maps from acquired (GT), standard-grid, and off-grid echoes, with absolute-error maps shown over $[0, 250]$ ms.

Table 2: Component analysis over ten subjects (mean±SD). PSNR averages five targets per FA; map metrics use the fixed acquired-reference mask.

	Mean echo-image quality		Echo-fitted T_2^* map quality	
Configuration	FA 4° PSNR↑	FA 16° PSNR↑	MAE↓	R^2 ↑
Pointwise MLP	34.74 ± 0.86	35.21 ± 1.09	15.04 ± 1.79	0.214 ± 0.067
Constant-Rate	34.26 ± 0.80	35.00 ± 1.11	14.83 ± 1.76	0.249 ± 0.068
Unconstrained Decoder	35.08 ± 0.97	$\mathbf{35.54 \pm 1.17}$	14.58 ± 1.79	0.255 ± 0.048
w/o Boundary Alignment	35.15 ± 0.96	35.47 ± 1.15	14.16 ± 1.57	0.312 ± 0.080
Full Model	$\mathbf{35.21 \pm 0.96}$	35.52 ± 1.15	$\mathbf{14.13 \pm 1.61}$	$\mathbf{0.318 \pm 0.088}$

4.4 What Do the Endpoint-Conditioned Trajectory Network and Physics-Constrained Decoder Contribute?

Table 2 compares four controls: the parameter-count-matched Pointwise MLP replaces the Endpoint-Conditioned Trajectory Network with direct acquisition-parameter-to-latent conditioning, retaining the encoder, multiplicative modulation, Physics-Constrained Decoder, losses, and training budget; Constant-Rate omits the Endpoint-Conditioned Trajectory Network and predicts one TE-independent positive rate per FA; Unconstrained Decoder replaces the Physics-Constrained Decoder with a direct image MLP; and w/o Boundary Alignment removes only the boundary objective.

Relative to the Full Model, Pointwise MLP lowers PSNR by $0.47/0.31$ dB and worsens map MAE/R^2 from $14.13/0.318$ to $15.04/0.214$, supporting endpoint-conditioned trajectory modeling over direct acquisition-parameter conditioning. Constant-Rate lowers PSNR by $0.95/0.52$ dB and reaches a T_2^* MAE of 14.83 ms and $R^2 = 0.249$, showing that a single mono-exponential rate per FA is less effective for this task.

The Unconstrained Decoder achieves competitive image fidelity and even slightly higher FA $16°$ PSNR than the Full Model (35.54 vs. 35.52 dB), yet worsens T_2^* MAE/R^2 from $14.13/0.318$ to $14.58/0.255$. The paired MAE difference is 0.45 ms (95% bootstrap CI, 0.24–0.68 ms; exact Wilcoxon $p = 0.0039$), with nine of ten subjects favoring the Full Model. Thus, competitive individual-echo fidelity alone does not guarantee preservation of the cross-TE relationships required for quantitative fitting. Boundary alignment has a smaller effect, changing PSNR by at most 0.06 dB and improving map MAE/R^2 from $14.16/0.312$ to $14.13/0.318$.

Physics-Constrained Decoder behavior beyond the supervised range. Figure 5 compares both decoders over dense TE queries from $t = 0.2$ to 1.4 (TE $= 3.86$–34.40 ms). Within the supervised range ($t \leq 1$), both decoders broadly follow the observed decay trend, consistent with their comparable echo-image PSNR. Beyond the supervised range, the Unconstrained Decoder flattens, while positive-rate integration preserves a smooth nonincreasing trajectory.

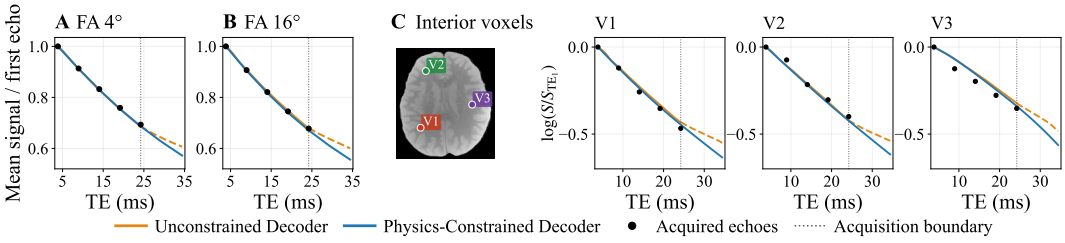


Figure 5: Effect of the Physics-Constrained Decoder on signal evolution. (A–B) Foreground-mean signals normalized at TE $= 3.86$ ms within the common foreground mask for FA $4°$ and $16°$ across ten subjects. Black markers denote acquired echoes, and the vertical dotted line marks the end of the supervised range ($t = 1$). (C) Prediction-error-blind interior voxels selected from acquired-reference T_2^* percentiles show the same decoder-dependent behavior in their FA $16°$ log-relative trajectories.

Because no acquired targets are available beyond $t = 1$, the beyond-range curves characterize decoder-induced continuation rather than extrapolation accuracy.

5 Discussion and Conclusion

We complete long-TR echo trains from retained short-TR echoes using continuous feature modulation and positive-rate integration, supporting flexible TE queries with nonnegative, nonincreasing signals. At comparable image fidelity, the Physics-Constrained Decoder improves acquired-reference T_2^* map agreement over the Unconstrained Decoder. Evidence is limited to one private healthy-participant cohort, scanner, protocol, subject split, and training seed; paired uncertainty is conditional on the trained models. Agreement with acquired-echo maps from one estimator establishes neither absolute tissue-parameter accuracy nor clinical utility. Regional predictions show reduced between-subject dispersion, leaving preservation of individual variation unresolved.

AI USE STATEMENT

Generative AI tools were used for English-language editing and translation; drafting and restructuring portions of the manuscript; literature search, summarization, and related-work identification; reference and formatting checks; internal consistency checks against saved evaluation outputs; review of claim calibration and result presentation; figure-layout and legibility adjustments; and small utility scripts for manuscript and figure preparation and verification.

Generative AI was not used to acquire or generate the research data, formulate the underlying model, run model training or inference, or produce the primary quantitative results. Retained citations were checked against source publications, numerical statements against saved evaluation outputs, and figure changes against the underlying results. The authors determined the study design, model formulation, experiments, quantitative analyses, scientific interpretation, and final claims, and take responsibility for the complete content of the submission, including AI-assisted text and artifacts.

ETHICS STATEMENT

This study uses a private brain MRI dataset acquired from 42 healthy participants. The data are not publicly available. The synthesized images and parameter maps are evaluated for research purposes and are not intended for clinical decision-making in this study.

REPRODUCIBILITY STATEMENT

The model architecture and training objectives are specified in Section 3, and the completion task and evaluation design are summarized in Section 4. Appendix B documents the subject split, acquisition protocol, normalization, optimization settings, fixed final-checkpoint rule, comparison implementations, and result provenance; it also specifies numerical stabilization and signal integration. Appendix C defines the T_2^* estimator, masks, invalid-value handling, and map-level metrics. Appendix A provides per-echo results and dense-query visualizations, while Appendix D reports subject-level, tissue-level, and regional quantitative summaries. These protocol and processing details are provided to support reproducibility. Code will be released upon acceptance.

REFERENCES

John Ashburner and Karl J. Friston. Unified segmentation. *NeuroImage*, 26(3):839–851, 2005. doi: 10.1016/j.neuroimage.2005.02.018.

J. Martin Bland and Douglas G. Altman. Statistical methods for assessing agreement between two methods of clinical measurement. *The Lancet*, 327(8476):307–310, 1986. doi: 10.1016/S0140-6736(86)90837-8.

Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems*, volume 31, 2018. URL https://papers.nips.cc/paper/2018/hash/69386f6bb1dfed68692a24c8686939b9-Abstract.html.

Onat Dalmaz, Mahmut Yurt, and Tolga Çukur. ResViT: Residual vision transformers for multimodal medical image synthesis. *IEEE Transactions on Medical Imaging*, 41(10):2598–2614, 2022. doi: 10.1109/TMI.2022.3167808.

Salman U. H. Dar, Mahmut Yurt, Levent Karacan, Aykut Erdem, Erkut Erdem, and Tolga Çukur. Image synthesis in multi-contrast MRI with conditional generative adversarial networks. *IEEE Transactions on Medical Imaging*, 38(10):2375–2388, 2019. doi: 10.1109/TMI.2019.2901750. URL https://pubmed.ncbi.nlm.nih.gov/30835216/.

Sanuwani Dayarathna, Himashi Peiris, Kh Tohidul Islam, Tien-Tsin Wong, and Zhaolin Chen. D2Diff: A dual-domain diffusion model for accurate multi-contrast MRI synthesis. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2025*, volume 15961, pp. 131–141, 2025. doi: 10.1007/978-3-032-04937-7_13. URL https://papers.miccai.org/miccai-2025/0207-Paper2110.html.

Sean C. L. Deoni, Brian K. Rutt, and Terry M. Peters. Rapid combined T1 and T2 mapping using gradient recalled acquisition in the steady state. *Magnetic Resonance in Medicine*, 49(3):515–526, 2003. doi: 10.1002/mrm.10407.

Hákon Gudbjartsson and Samuel Patz. The Rician distribution of noisy MRI data. *Magnetic Resonance in Medicine*, 34(6):910–914, 1995. doi: 10.1002/mrm.1910340618.

Amir Heydari, Abbas Ahmadi, Tae Hyung Kim, and Berkin Bilgic. Joint MAPLE: Accelerated joint T1 and T2* mapping with scan-specific self-supervised networks. *Magnetic Resonance in Medicine*, 91(6):2294–2309, 2024. doi: 10.1002/mrm.29989. URL https://pubmed.ncbi.nlm.nih.gov/38181183/.

George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021. doi: 10.1038/s42254-021-00314-5.

Yawen Liu, Hongxia Yin, Zuofeng Zheng, Wenjuan Liu, Tingting Zhang, Linkun Cai, Haijun Niu, Han Lv, Zhenghan Yang, Zhenchang Wang, and Pengling Ren. Two-minute deep learning–powered brain quantitative mapping: Accelerating clinical imaging with synthetic magnetic resonance imaging. *JMIR Medical Informatics*, 14:e79389, 2026. doi: 10.2196/79389.

Dan Ma, Vikas Gulani, Nicole Seiberlich, Kecheng Liu, Jeffrey L. Sunshine, Jeffrey L. Duerk, and Mark A. Griswold. Magnetic resonance fingerprinting. *Nature*, 495(7440):187–192, 2013. doi: 10.1038/nature11971.

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *Computer Vision – ECCV 2020*, volume 12346, pp. 405–421, 2020. doi: 10.1007/978-3-030-58452-8_24.

Dinor Nagar, Sahar Ifrah, Alex Finkelstein, Nikita Vladimirov, Moritz Zaiss, and Or Perlman. Multi-contrast generation and quantitative MRI using a transformer-based framework with RF excitation embeddings. *Communications Biology*, 9:102, 2026. doi: 10.1038/s42003-025-09371-3.

Niklas Nolte, Ouail Kitouni, and Mike Williams. Expressive monotonic neural networks. In *International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=w2P7fMy_RH.

Shihan Qiu, Sen Ma, Lixia Wang, Yuhua Chen, Zhaoyang Fan, Franklin G. Moser, Marcel Maya, Pascal Sati, Nancy L. Sicotte, Anthony G. Christodoulou, Yibin Xie, and Debiao Li. Direct synthesis of multi-contrast brain MR images from MR multitasking spatial factors using deep learning. *Magnetic Resonance in Medicine*, 90(4):1672–1681, 2023. doi: 10.1002/mrm.29715.

Vincent Sitzmann, Julien N. P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *Advances in Neural Information Processing Systems*, volume 33, pp. 7462–7473, 2020. URL https://papers.nips.cc/paper/2020/hash/53c04118df112c13a8c34b38343b9c10-Abstract.html.

Paul S. Tofts (ed.). *Quantitative MRI of the Brain: Measuring Changes Caused by Disease*. John Wiley & Sons, 2003. doi: 10.1002/0470869526.

Nathalie Tzourio-Mazoyer, Brigitte Landeau, Dimitri Papathanassiou, Fabrice Crivello, Olivier Etard, Nicolas Delcroix, Bernard Mazoyer, and Marc Joliot. Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *NeuroImage*, 15(1):273–289, 2002. doi: 10.1006/nimg.2001.0978.

Ke Wang, Mariya Doneva, Jakob Meineke, Thomas Amthor, Ekin Karasan, Fei Tan, Jonathan I. Tamir, Stella X. Yu, and Michael Lustig. High-fidelity direct contrast synthesis from magnetic resonance fingerprinting. *Magnetic Resonance in Medicine*, 90(5):2116–2129, 2023. doi: 10.1002/mrm.29766.

Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004. doi: 10.1109/TIP.2003.819861.

Antoine Wehenkel and Gilles Louppe. Unconstrained monotonic neural networks. In *Advances in Neural Information Processing Systems*, volume 32, 2019. URL https://proceedings.neurips.cc/paper/2019/hash/2a084e55c87b1ebcdaad1f62fdbbac8e-Abstract.html.

Yicheng Wu, Tao Song, Zhonghua Wu, Jin Ye, Zongyuan Ge, Wenjia Bai, Zhaolin Chen, and Jianfei Cai. Virtual full-stack scanning of brain MRI via imputing any quantised code. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 21026–21035, 2026.

Honglin Xiong, Yu Fang, Kaicong Sun, Yulin Wang, Xiaopeng Zong, Weijun Zhang, and Qian Wang. Contrast representation learning from imaging parameters for magnetic resonance image synthesis. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2024*, volume 15007, pp. 181–190, 2024. doi: 10.1007/978-3-031-72104-5_18.

Honglin Xiong, Yulin Wang, Zhenrong Shen, Kaicong Sun, Yu Fang, Yan Chen, Dinggang Shen, and Qian Wang. Learning contrast and content representations for synthesizing magnetic resonance image of arbitrary contrast. *Medical Image Analysis*, 104:103635, 2025. doi: 10.1016/j.media.2025.103635.

Yongquan Ye, Jingyuan Lyu, Yichen Hu, Zhongqi Zhang, Jian Xu, and Weiguo Zhang. MULTIparametric MR imaging with fLEXible design (MULTIPLEX). *Magnetic Resonance in Medicine*, 87(2):658–673, 2022. doi: 10.1002/mrm.28999.

Xinzhe Zhang, Junjie Liang, Peng Cao, Jinzhu Yang, and Osmar R. Zaiane. Structure-aware MRI translation: Multi-modal latent diffusion model with arbitrary missing modalities. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2025*, volume 15967, pp. 508–518, 2025. doi: 10.1007/978-3-032-04984-1_49.

A PER-ECHO RESULTS FOR TABLE 1

Tables 3 and 4 expand the image-quality columns of Table 1. Echo indices run from 1 to 5 within each FA, corresponding to TE $3.86, 8.95, 14.04, 19.13, 24.22$ ms. Each cell reports the mean and population SD across the same ten evaluation subjects; the FA-level mean is the average of the five echo means. The SD in Table 1 is computed after averaging echoes within each subject and therefore is not the average of the five SDs. The joint T_2^* endpoints are reported once in Table 1, since they use both complete echo trains.

Table 3: Per-echo image fidelity at FA 4° over ten subjects. Higher PSNR (dB) and SSIM are better.

Method	Metric	Echo 1	Echo 2	Echo 3	Echo 4	Echo 5
	TE (ms)	3.86	8.95	14.04	19.13	24.22
ResViT	PSNR	36.58 ± 1.00	34.66 ± 1.00	32.17 ± 1.19	30.56 ± 1.21	29.50 ± 1.28
	SSIM	0.9694 ± 0.0031	0.9556 ± 0.0052	0.9354 ± 0.0086	0.9192 ± 0.0119	0.9056 ± 0.0142
D2Diff	PSNR	42.22 ± 0.49	37.34 ± 0.98	33.71 ± 1.24	31.48 ± 1.28	30.11 ± 1.32
	SSIM	0.9781 ± 0.0025	0.9675 ± 0.0038	0.9513 ± 0.0070	0.9365 ± 0.0094	0.9215 ± 0.0118
CodeBrain	PSNR	35.12 ± 0.70	33.62 ± 0.84	31.61 ± 1.02	30.07 ± 1.14	28.97 ± 1.21
	SSIM	0.9623 ± 0.0029	0.9475 ± 0.0053	0.9284 ± 0.0087	0.9101 ± 0.0121	0.8947 ± 0.0146
CCR	PSNR	41.38 ± 0.50	36.87 ± 0.99	33.53 ± 1.13	31.42 ± 1.19	30.07 ± 1.21
	SSIM	0.9767 ± 0.0026	0.9645 ± 0.0039	0.9472 ± 0.0069	0.9308 ± 0.0100	0.9161 ± 0.0123
Ours	PSNR	42.10 ± 0.49	37.38 ± 0.99	34.04 ± 1.13	31.94 ± 1.19	30.56 ± 1.24
	SSIM	0.9779 ± 0.0025	0.9668 ± 0.0038	0.9518 ± 0.0066	0.9370 ± 0.0094	0.9235 ± 0.0118

Table 4: Per-echo image fidelity at FA 16° over ten subjects. Higher PSNR (dB) and SSIM are better.

Method	Metric	Echo 1	Echo 2	Echo 3	Echo 4	Echo 5
	TE (ms)	3.86	8.95	14.04	19.13	24.22
ResViT	PSNR	36.54 ± 1.37	34.10 ± 1.28	31.93 ± 1.39	31.29 ± 1.36	28.69 ± 1.44
	SSIM	0.9829 ± 0.0020	0.9651 ± 0.0052	0.9432 ± 0.0094	0.9253 ± 0.0117	0.8945 ± 0.0160
D2Diff	PSNR	42.84 ± 0.50	36.63 ± 1.39	33.74 ± 1.37	32.08 ± 1.42	29.19 ± 1.39
	SSIM	0.9892 ± 0.0008	0.9746 ± 0.0037	0.9555 ± 0.0072	0.9378 ± 0.0100	0.9102 ± 0.0126
CodeBrain	PSNR	34.82 ± 0.79	33.48 ± 0.99	32.02 ± 1.14	30.99 ± 1.26	28.44 ± 1.35
	SSIM	0.9757 ± 0.0015	0.9581 ± 0.0047	0.9375 ± 0.0084	0.9178 ± 0.0118	0.8863 ± 0.0153
CCR	PSNR	41.23 ± 0.55	36.10 ± 1.32	33.76 ± 1.34	32.42 ± 1.36	29.76 ± 1.37
	SSIM	0.9867 ± 0.0008	0.9712 ± 0.0039	0.9528 ± 0.0073	0.9354 ± 0.0104	0.9095 ± 0.0139
Ours	PSNR	42.35 ± 0.49	36.59 ± 1.34	34.09 ± 1.36	32.69 ± 1.41	31.87 ± 1.42
	SSIM	0.9889 ± 0.0008	0.9740 ± 0.0037	0.9573 ± 0.0067	0.9414 ± 0.0096	0.9271 ± 0.0122

Complete and continuously queried echo trains. Figure 6 provides a qualitative counterpart to the per-echo measurements above by showing the complete acquired and synthesized echo trains under both flip angles.

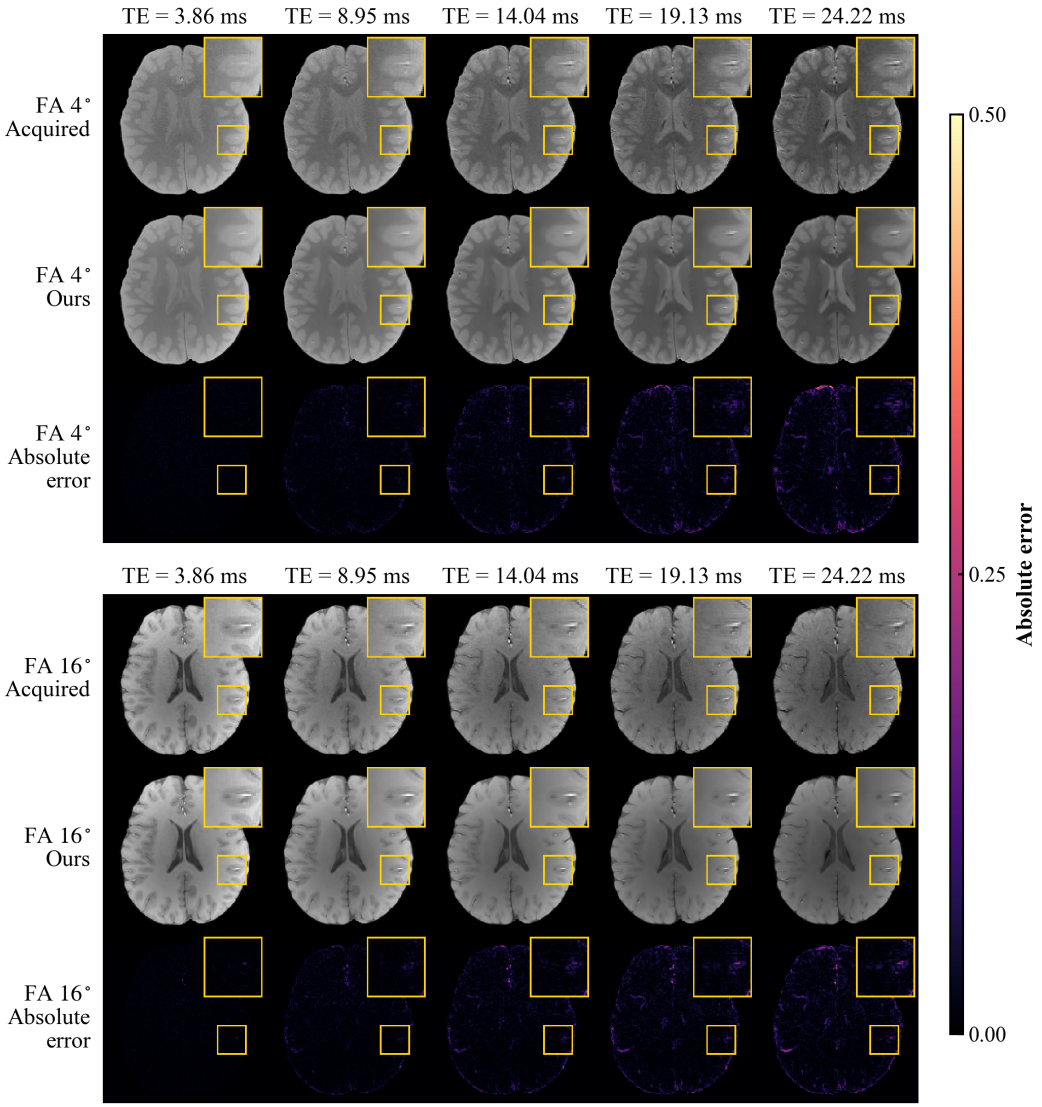


Figure 6: Complete synthesized target echo trains for subject 06 (slice 6). For each flip angle, the columns correspond to the five acquired TEs and the rows show the acquired echoes, our predictions, and normalized absolute errors under a common $[0, 0.5]$ display range. Yellow boxes identify a fixed anatomical region, with matched magnifications shown in the upper-right corner. The visualization complements Tables 3 and 4 by showing how anatomical detail and residual error evolve across the entire echo train.

B Evaluation Protocol, Comparisons, and Implementation

This appendix documents the acquisition protocol, evaluation scope, implementation details, saved-map fitting conventions, and result provenance referenced from the main text.

Data and acquisition. We acquired this private brain MRI dataset from 42 healthy participants using a United Imaging uMR 890 3T scanner with a 64-channel head coil. The dataset is private and is not publicly available. The current retrospective experiments use 32 training subjects and the ten evaluation subjects 01, 02, 04–11. For each FA, the retained source belongs to the TR1 $= 7.97$ ms block and the five targets belong to the TR2 $= 29.23$ ms block; equality of a source TE and a target TE therefore does not imply equal signal amplitude. Consistent with the main text, the two target groups correspond to FA $= 4^\circ$ and 16°, and each contains the acquired TEs $\{3.86, 8.95, 14.04, 19.13, 24.22\}$ ms.

Under equal numbers of repetitions and unchanged overheads, the retained fraction of the two-block repetition-time budget is $7.97/(7.97 + 29.23) = 0.214$. This corresponds to a nominal 78.6% reduction of that component; prospective wall-clock acceleration was not measured.

Selection and scope of evaluation. All reported metrics are computed after inference on the full evaluation volumes. All models are trained for the fixed 1000-epoch budget, and all reported evaluations use the final-epoch checkpoint. The ten evaluation subjects are not used for parameter updates, checkpoint selection, early stopping, or learning-rate scheduling. In the 14.04-ms held-out experiment, the 14.04-ms target image of each FA group is completely excluded from training: it is not read by the data loader and does not enter reconstruction loss, tissue-basis construction, or any other training objective. The 24.22-ms boundary experiment applies the same complete exclusion to the 24.22-ms target image of each FA group. In both experiments, the withheld images are used only for final evaluation; only the corresponding numerical acquisition coordinates remain available to the conditioning network. Consequently, both protocols test acquisition-coordinate querying without direct image supervision at the specified TE; neither changes the evaluation subjects.

External comparisons. The comparison set comprises ResViT, D2Diff, CodeBrain, and CCR. All methods use the same 32/10 subject split and receive the same single retained source echo per FA. The selected experiments use target-specific models for ResViT/D2Diff and FA-specific one-to-many models for CodeBrain/CCR; our network shares weights across both FAs. Each comparison method was obtained from the corresponding authors' official released code and retrained from scratch on the same private dataset and 32/10 subject split used by our model; no author-provided pretrained weights were used. We retained each method's official architecture and method-specific training objective while harmonizing the retained-source inputs, target echoes, fixed training budget, and evaluation data. Inference outputs from all methods were then processed by the common evaluation pipeline described below. For the downstream comparison, every method's synthesized echo train is processed by the same pooled stride-two ratio estimator and scored against the same acquired reference with the same fixed acquired-reference mask described in Appendix C.

Normalization and evaluation. The current training and inference use corrected-intensity-normalized volumes. Within each FA group, the retained source and all target echoes use the same fixed intensity scale, $c_{4^\circ} = 800$ or $c_{16^\circ} = 1000$, including the final target echo. For an acquired image $I^{(g)}_{k,\mathrm{raw}}$, preprocessing is

$$\widetilde{I}^{(g)}_k = \frac{\mathrm{clip}\Big(I^{(g)}_{k,\mathrm{raw}}, 0, c_g\Big)}{c_g}. \tag{14}$$

Thus both encoder inputs and supervised targets lie in $[0, 1]$, while the use of a common fixed scale within each FA group preserves the relative source-to-target and cross-TE signal amplitudes. Image PSNR/SSIM are computed in this normalized space with data range one. Synthesized echoes are likewise clipped to $[0, 1]$ before evaluation. Before T_2^* fitting, both acquired and synthesized normalized echoes are restored to the common signal scale as $I^{(g)}_{k,\mathrm{fit}} = c_g \widetilde{I}^{(g)}_k$. The direct-query signal analysis in Section 4.4 instead evaluates the network trajectory itself and does not construct curves by interpolating the five acquired samples. Echo metrics are averaged over the five targets

within each FA and subject, followed by the mean and SD across subjects. In Section 4.3, the TE-specific rows evaluate only the 14.04-ms or the 24.22-ms target of each FA group. Exported echo SDs use $\mathrm{ddof} = 0$.

Model and training details. The image encoder has 32 residual blocks, 256 channels, residual scaling 0.1, and no batch normalization. Parameter encoding and the trajectory generator use two and three 256-unit hidden layers, respectively. The amplitude and rate decoders have four and three 256-unit hidden layers. Hidden activations are ReLU; signed linear latent outputs precede the specified signal softplus operations. The amplitude uses softplus with $\beta = 10$, and the rate head is initialized to $1/30\ \mathrm{ms}^{-1}$.

Training uses Adam, learning rate 10^{-4} halved at epochs 200/400/600/800, batch size 8, 64×64 random patches, gradient-norm clipping at 5, seed 2026, and a 1000-epoch budget. Evaluation uses 180×180 center crops. Data are repeated twice without augmentation, with four data-loader workers. Our model uses weights $(1, 0.1)$ for echo and boundary losses; the ablation without boundary alignment changes the boundary weight to zero.

Result provenance. The Full Model and every comparison, component analysis, and held-out-TE configuration use the checkpoint from the final epoch of the same pre-specified training budget. No reported checkpoint is selected using performance on the ten evaluation subjects.

All saved maps are evaluated directly against the current acquired-echo reference. Echo SDs retain the exported $\mathrm{ddof} = 0$ convention; map SDs use the sample convention $\mathrm{ddof} = 1$.

Constant-Rate control. This simple baseline uses the same retained source image and source acquisition descriptor to predict a nonnegative reference amplitude $A_g(x)$ and one TE-independent positive rate $r_g(x)$ per FA. It omits the Endpoint-Conditioned Trajectory Network and generates every target echo as

$$\widehat{S}_g(x, \mathrm{TE}) = A_g(x) \exp[-r_g(x)(\mathrm{TE} - \mathrm{TE}_1)]\,. \tag{15}$$

It receives only image supervision. Its evaluation yields FA 4° PSNR/SSIM of $34.26 \pm 0.80/0.946 \pm 0.007$ and FA 16° PSNR/SSIM of $35.00 \pm 1.11/0.955 \pm 0.007$. Echo-fitted T_2^* evaluation gives an MAE of 14.83 ± 1.76 ms and map-agreement R^2 of 0.249 ± 0.068.

Table 5: Per-echo image results for the Constant-Rate control over ten subjects (mean±SD). MAE is computed in normalized image intensity. Echo indices before and after the slash correspond to FA 4° and FA 16°, respectively.

	FA 4°			FA 16°		
Echo	PSNR	SSIM	MAE	PSNR	SSIM	MAE
1 / 7	40.66 ± 0.39	0.978 ± 0.003	0.00460 ± 0.00030	41.67 ± 0.45	0.989 ± 0.001	0.00404 ± 0.00029
2 / 8	36.25 ± 0.77	0.964 ± 0.004	0.00687 ± 0.00058	35.98 ± 1.25	0.973 ± 0.004	0.00661 ± 0.00085
3 / 9	33.22 ± 0.91	0.945 ± 0.007	0.00901 ± 0.00087	33.65 ± 1.31	0.954 ± 0.007	0.00822 ± 0.00119
4 / 10	31.26 ± 1.02	0.929 ± 0.010	0.01087 ± 0.00117	32.30 ± 1.40	0.937 ± 0.011	0.00951 ± 0.00152
5 / 11	29.92 ± 1.09	0.914 ± 0.013	0.01247 ± 0.00145	31.40 ± 1.46	0.921 ± 0.014	0.01059 ± 0.00180
Mean	34.26 ± 0.80	0.946 ± 0.007	0.00876 ± 0.00086	35.00 ± 1.11	0.955 ± 0.007	0.00780 ± 0.00110

Pointwise MLP control. This control replaces $G_\theta([\mathbf{c}_0, \mathbf{c}_1, t])$ with a direct MLP mapping from acquisition parameters $\mathbf{p} = (\mathrm{FA}, \mathrm{TR}, \mathrm{TE})$ to a 256-dimensional modulation vector. The source and target TR values are 7.97 and 29.23 ms, respectively:

$$\mathbf{z}_{\mathrm{pw}}(\mathbf{p}) = P_\psi(\mathbf{p}), \qquad \mathbf{B}_{\mathrm{pw}}(x) = \frac{E_{\mathrm{img}}(I_{\mathrm{echo}_0})(x)}{d_\epsilon(\mathbf{z}_{\mathrm{pw}}(\mathbf{p}_{\mathrm{echo}_0}))}, \qquad \mathbf{F}_{\mathrm{pw}}(x, \mathbf{p}) = \mathbf{B}_{\mathrm{pw}}(x) \odot \mathbf{z}_{\mathrm{pw}}(\mathbf{p}). \tag{16}$$

Thus the change concerns acquisition conditioning; the source encoder, stabilized factorization, multiplicative modulation, amplitude/rate heads, and positive-rate integration are retained. Both models retain the acquisition encoder for the boundary objective and use the same loss weights, data split, preprocessing, and 1000-epoch optimization schedule specified above. Hidden layers use ReLU and latent outputs are linear. The total parameter counts differ by only 264 (0.00066%).

For held-out comparisons, the specified target is excluded from image reads and supervised losses for both models. Table 2 reports full-supervision results and Table 7 reports the held-out endpoints.

Table 6: Architecture and parameter counts for the Endpoint-Conditioned Trajectory Network and Pointwise MLP controls. Layer widths include input and output; total counts include the shared acquisition-encoder architecture.

Configuration	Conditioning network widths	Module parameters	Total parameters
Full Model	$513 \to 256 \to 256 \to 256 \to 256$	328,960	39,741,701
Pointwise MLP	$3 \to 345 \to 345 \to 345 \to 256$	328,696	39,741,437

Unconstrained Decoder control. The Unconstrained Decoder row in Table 2 is an architecture-controlled ablation of the Physics-Constrained Decoder. It retains the same image encoder, acquisition encoder, tissue factorization, Endpoint-Conditioned Trajectory Network, training data, objectives, and optimization schedule as the Full Model. It replaces the physics-constrained amplitude–rate formulation and positive-rate exponential integration with a direct image MLP having four 256-unit ReLU hidden layers and a linear one-channel output. The control contains $39{,}544{,}068$ trainable parameters, compared with $39{,}741{,}701$ for the Full Model, a difference of approximately 0.50%. Its FA $4°$/FA$16°$ mean PSNRs are 35.08 ± 0.97 and 35.54 ± 1.17 dB, and the corresponding SSIMs are 0.9496 ± 0.0068 and 0.9573 ± 0.0065. Under the fixed acquired-reference mask, its echo-fitted T_2^* MAE is 14.58 ± 1.79 ms and map-agreement R^2 is 0.255 ± 0.048.

B.1 Complete Held-Out Echo Results

Table 7: Prediction at acquired TE coordinates withheld from image supervision (ten subjects; mean±SD). "Full Model: full" denotes the reference model supervised at all five target echoes; all other rows withhold the stated target. Image metrics are computed only at the indicated query; T_2^* metrics are computed from the completed five-echo trains jointly across both FAs using the same fixed acquired-reference mask as Table 1. All rows use the final-epoch checkpoint from their respective fixed-budget training runs.

		FA $4°$		FA $16°$		Joint T_2^* fit	
Query	Configuration	PSNR↑	SSIM↑	PSNR↑	SSIM↑	MAE↓	R^2 ↑
14.04 ms	Full Model: full	34.04 ± 1.13	0.9518 ± 0.0066	34.09 ± 1.36	0.9573 ± 0.0067	14.13 ± 1.61	0.318 ± 0.088
14.04 ms	Full Model: held out	34.01 ± 1.12	0.9518 ± 0.0065	34.04 ± 1.33	0.9569 ± 0.0068	14.17 ± 1.58	0.306 ± 0.083
14.04 ms	Pointwise: held out	33.49 ± 1.02	0.9461 ± 0.0073	33.64 ± 1.26	0.9533 ± 0.0073	15.08 ± 1.75	0.207 ± 0.077
24.22 ms	Full Model: full	30.56 ± 1.24	0.9235 ± 0.0118	31.87 ± 1.42	0.9271 ± 0.0122	14.13 ± 1.61	0.318 ± 0.088
24.22 ms	Full Model: held out	30.52 ± 1.26	0.9236 ± 0.0119	31.84 ± 1.43	0.9269 ± 0.0124	14.28 ± 1.59	0.295 ± 0.080
24.22 ms	Pointwise: held out	30.01 ± 1.17	0.9127 ± 0.0131	31.46 ± 1.46	0.9206 ± 0.0131	15.02 ± 1.74	0.223 ± 0.067

Numerical query interface. The curves in Figure 5 are constructed from direct, densely sampled model queries rather than interpolation of the five acquired echoes. Each subject–decoder curve is normalized by its own value at TE $= 3.86$ ms before aggregation.

Equation (7) evaluates the reference amplitude at $t = 0.2$, while Equations (9) and (18) use the corresponding physical reference TE of 3.86 ms. All reported continuous-trajectory analyses use a fixed, ordered query grid originating at this reference echo. The same in-range grid is used across subjects and decoder comparisons, and cumulative trapezoidal integration is evaluated under this declared protocol. Consequently, every plotted coordinate is a reproducible network query on a common numerical grid, including the unacquired intermediate TEs.

B.2 Quadrature-Grid Consistency

Because Eq. (18) is evaluated discretely, we test whether the prediction at a common target TE changes materially when intermediate integration nodes are added. Without retraining, the same final-epoch Full Model checkpoint used for Table 1 is evaluated on three ordered grids over $t \in [0.2, 1.0]$: the five acquired coordinates $\{0.2, 0.4, 0.6, 0.8, 1.0\}$; a 121-point uniform grid; and a fixed 17-point nonuniform grid $\{0.2, 0.225, 0.27, 0.33, 0.4, 0.49, 0.56, 0.6, 0.69, 0.76, 0.8, 0.86, 0.91, 0.95, 0.975, 0.99, 1.0\}$. Each grid contains all five acquired coordinates exactly. We compare predictions only at these common coordinates, using full-volume signals on the normalized image scale. We then refit the joint magnitude-domain T_2^* map and score it on the same fixed acquired-reference mask defined below.

The five-point baseline therefore reproduces the Table 1 Full Model MAE up to unrounded numerical precision.

Table 8: Quadrature-grid consistency across ten subjects. Image MAE and PSNR compare common-TE predictions with those from the original five-point grid. Map ΔMAE compares the fitted T_2^* maps directly, whereas the last two columns report MAE against the acquired-echo reference and its change from the five-point result. Values are mean±SD across subjects.

Grid	Nodes	Image MAE↓	Image PSNR↑ (dB)	Map ΔMAE↓ (ms)	Reference MAE↓ (Δ) (ms)
Acquired	5	–	–	0.000	14.125 ± 1.611 (0.000)
Uniform	121	$(0.76 \pm 0.11) \times 10^{-4}$	70.56 ± 1.61	0.182 ± 0.039	14.173 ± 1.612 (+0.047)
Nonuniform	17	$(0.68 \pm 0.10) \times 10^{-4}$	71.45 ± 1.58	0.162 ± 0.035	14.167 ± 1.612 (+0.042)

The common-TE image differences are below 10^{-4} in normalized MAE and exceed 70 dB in cross-grid PSNR. Changing the integration grid alters the fitted maps by only 0.16–0.18 ms on average and changes reference-map MAE by at most 0.047 ms. Thus, the discrete prediction is not mathematically grid invariant, but the evaluated image and quantitative endpoints are numerically stable across the tested grids. All other experiments retain their declared fixed grids for exact reproducibility.

B.3 Numerical Feature Stabilization and Signal Integration

Stable feature factorization. In Equation (4), the component-wise denominator is

$$d_\epsilon(u) = \begin{cases} \max(u, \epsilon), & u \geq 0, \\ -\max(-u, \epsilon), & u < 0. \end{cases} \tag{17}$$

with $\epsilon = 10^{-2}$. This preserves the sign of nonzero latent components and bounds the denominator away from zero, with $d_\epsilon(0) = \epsilon$. At $t = t_0$ the ratio is unity except on channels with $|\mathbf{z}(t_0)| < \epsilon$, where the stabilizer keeps the ratio finite rather than exactly one.

Cumulative signal integration. For an ordered query set $\mathcal{Q} = \{t_j\}_{j=1}^M$ starting at the reference coordinate t_1, let $\widetilde{\mathrm{TE}}_j$ be the corresponding physical TEs and $r_j(x) = r(x, \widetilde{\mathrm{TE}}_j)$. The implementation uses cumulative trapezoidal integration:

$$H_1(x) = 0, \qquad H_j(x) = H_{j-1}(x) + \tfrac{1}{2}\big(r_{j-1}(x) + r_j(x)\big)(\widetilde{\mathrm{TE}}_j - \widetilde{\mathrm{TE}}_{j-1}), \quad j = 2, \ldots, M, \tag{18}$$

$$\widehat{S}_j(x) = A(x) e^{-H_j(x)}, \qquad j = 1, \ldots, M. \tag{19}$$

Why the constraint holds. For $A \geq 0$ and $r \geq 0$, Eq. (9) gives $\widehat{S} \geq 0$ and $\partial_{\mathrm{TE}} \widehat{S} = -r\widehat{S} \leq 0$. On an ordered numerical grid, every trapezoidal increment in Eq. (18) is nonnegative. Hence $H_{j+1} \geq H_j$ and $Ae^{-H_{j+1}} \leq Ae^{-H_j}$. This establishes the same ordering for the numerical synthesis rule. For a positive differentiable decreasing signal, the corresponding rate is $-\partial_{\mathrm{TE}} \log S$; a constant rate yields mono-exponential decay. The constraint specifies a signal family, while reconstruction learns which member fits each input.

C Fitting and Signal-Diagnostic Conventions

Joint quantitative estimator. The scanner-side MULTIPLEX T_2^* fitting implementation was not available to us. We therefore used an independent research implementation of the estimator described in the published MULTIPLEX formulation (Ye et al., 2022); the magnitude-domain estimator used here is specified in Eq. (20). The same fitting code processes the synthesized echo trains from ResViT, D2Diff, CodeBrain, CCR, and our model, as well as the acquired reference echoes. All methods are subsequently evaluated using the same fixed acquired-reference mask and scoring procedure. All ten target echoes, excluding the source images, are restored to their FA-specific intensity scale and jointly fitted. For finite positive same-stride pairs, the implemented magnitude-domain

pooled ratio is

$$q(x) = \frac{\sum_g \sum_{j=1}^{3} S_g(x, \mathrm{TE}_j) S_g(x, \mathrm{TE}_{j+2})}{\sum_g \sum_{j=1}^{3} S_g(x, \mathrm{TE}_j)^2}, \qquad \widehat{T}_2^*(x) = -\frac{10.18\ \mathrm{ms}}{\log q(x)}. \tag{20}$$

The same estimator is applied to the generated and acquired trains. This is a post-synthesis computation, not the inverse of a network rate or an objective used during synthesis training.

Fixed-reference evaluation of saved maps. For each subject s, the initial fitting support is the union of voxels whose normalized intensity exceeds 10^{-6} in at least one acquired target echo. Within that support, finite same-polarity echo pairs with both signals above 10^{-6} contribute to Eq. (20). Let E_s denote the acquired-reference estimated mask: $E_s(v) = 1$ when at least one such pair is valid, the pooled denominator is nonzero, and the resulting ratio satisfies $10^{-6} < q(v) < 1$, yielding a finite positive raw T_2^* estimate. Let T_s be the reference map computed once from the acquired echo trains using this common estimator. We freeze $\mathcal{M}_s = \{v : E_s(v) = 1,\ T_s(v) > 0,\ T_s(v) \text{ is finite}\}$ and reuse exactly the same voxel indices for every method. This is a reference-derived estimability mask, not an anatomical segmentation: it uses no erosion, connected-component operation, prediction-derived criterion, relative-decay threshold, or fit-R^2 threshold. Evaluation uses the full saved three-dimensional volume after checking image shape and affine alignment. No prediction-reliability or decay-fit-R^2 threshold changes this mask. Nonfinite or nonpositive predictions are assigned the existing failure value zero and retained in the evaluation; their fraction is recorded separately. At this scoring stage, maps already produced by the common fitting implementation are loaded without further refitting; finite positive estimates above 500 ms have been clipped to 500 ms, and this truncation is retained.

$$\mathrm{MAE}_s = \frac{1}{|\mathcal{M}_s|} \sum_{v \in \mathcal{M}_s} |\widehat{T}_s(v) - T_s(v)|, \qquad R_s^2 = 1 - \frac{\sum_{v \in \mathcal{M}_s} (\widehat{T}_s(v) - T_s(v))^2}{\sum_{v \in \mathcal{M}_s} (T_s(v) - \overline{T}_s)^2}, \tag{21}$$

where $\overline{T}_s$ is the reference mean within $\mathcal{M}_s$. This R^2 measures map agreement, is not squared correlation, and can be negative. Each metric is computed per subject before averaging over the ten subjects; map-metric SDs use denominator nine. The fixed-mask result for the current model is 14.13 ± 1.61 ms and 0.318 ± 0.088. Reference maps, masks, per-subject metrics, and mask hashes accompany the evaluation export.

Signal diagnostics. For a fixed acquired-data mask M, J query times and a declared tolerance ϵ_{num}, define

$$V_{\mathrm{inc}} = \frac{\sum_{x \in M} \sum_{j=1}^{J-1} \mathbf{1}[\widehat{S}(x, \mathrm{TE}_{j+1}) > \widehat{S}(x, \mathrm{TE}_j) + \epsilon_{\mathrm{num}}]}{|M|(J-1)}, \tag{22}$$

$$V_{\mathrm{neg}} = \frac{\sum_{x \in M} \sum_{j=1}^{J} \mathbf{1}[\widehat{S}(x, \mathrm{TE}_j) < 0]}{|M|J}. \tag{23}$$

Nonfinite outputs are separate failures. Observed-range and beyond-range pairs are reported separately, including the first transition beyond the observed boundary. The fitting pipeline separately reports a decay-violation voxel count, which is not the raw adjacent-pair fraction defined above. For the Physics-Constrained Decoder, nonnegativity and monotonicity hold before image-range clipping: the amplitude mapping gives $A \geq 0$, the positive-rate parameterization makes every ordered quadrature increment nonnegative, and therefore $\widehat{S}(x, \mathrm{TE}_{j+1}) \leq \widehat{S}(x, \mathrm{TE}_j)$ directly at the decoder output. Accordingly, absent nonfinite numerical failures, the structural model gives $V_{\mathrm{neg}} = V_{\mathrm{inc}} = 0$ on any ordered query grid. The Full Model and the model without boundary alignment are consistent with this guarantee in all ten exported cases. These diagnostics complement the architectural guarantee and are not used for checkpoint selection.

Uncertainty. Subject identities are the sampling units. The reported dispersion describes ten subjects from a single training seed. Tissue and region scatter points are descriptive repeated measurements; inference across independent cohorts requires a separate validation study.

D Quantitative Evaluation Details

D.1 Subject-Level Map Agreement

Table 9 expands the map endpoints of the decoder and boundary comparisons. The same acquired-reference voxel mask is used for all three configurations in each subject. Mean and sample SD are computed across subjects, retaining failed predictions in the scoring mask as specified in Appendix C. For the prespecified decoder comparison, paired MAE differences are formed as Unconstrained Decoder minus Full Model, so a positive value favors the Full Model. The mean paired improvement is 0.454 ms (percentile paired-bootstrap 95% CI, 0.237–0.675 ms; 200,000 subject-level resamples, seed 2026). A two-sided exact Wilcoxon signed-rank test gives $W = 1$ and $p = 0.0039$. The corresponding paired effect sizes are Cohen's $d_z = 1.22$ and matched-pairs rank-biserial correlation $r_{\mathrm{rb}} = 0.964$; nine of ten subjects favor the Full Model. These subject-level MAE analyses quantify evaluation-cohort uncertainty conditional on the single training seed. We do not treat the aggregate R^2 difference as a second paired hypothesis test.

Table 9: Subject-level T_2^* map MAE (ms) used for the paired decoder analysis.

Subject	Full Model	No boundary	Unconstrained Decoder
01	13.07	13.31	13.45
02	13.62	13.60	14.33
04	15.18	15.32	15.75
05	14.03	14.09	14.61
06	12.71	12.71	12.88
07	18.09	18.06	19.15
08	13.88	13.93	14.03
09	13.88	13.77	14.05
10	12.41	12.62	13.31
11	14.37	14.18	14.24
Mean±SD	14.13 ± 1.61	14.16 ± 1.57	14.58 ± 1.79

D.2 Tissue and Regional Aggregation

GM and WM are defined from SPM tissue probabilities (Ashburner & Friston, 2005) thresholded at 0.80 after resampling to the map grid, followed by one six-connected binary erosion. Regional means use the bilateral AAL labels (Tzourio-Mazoyer et al., 2002) in Table 10. The saved common reliability mask is the intersection of the acquired-reference and prediction reliability masks. On each side, reliability requires a finite raw estimate in $[1, 500]$ ms, a fitted first-to-last relative decay of at least 0.03, and shared mono-exponential log-fit $R^2 \geq 0.80$. Each tissue or region is intersected with this method-specific common mask, and finite map values in $[0, 500]$ ms are retained. This produces a paired regional analysis conditional on reliable fitted values, whereas Table 1 retains the fixed, prediction-independent acquired-reference mask for all methods.

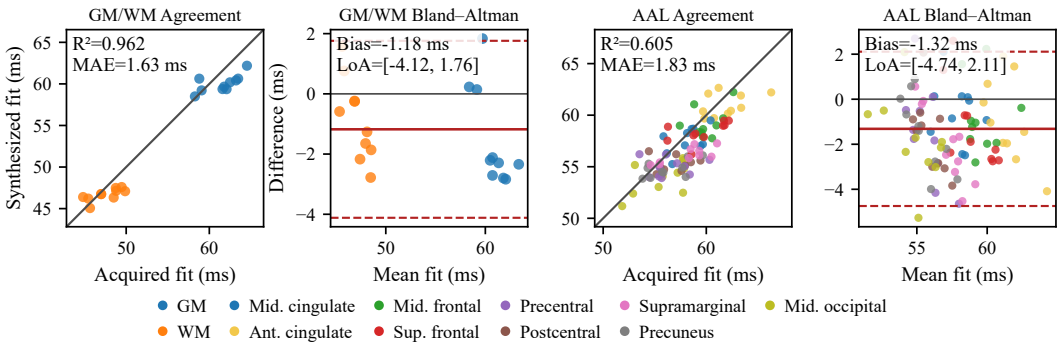


Figure 7: Tissue- and region-level agreement between T_2^* maps fitted from acquired and synthesized echo trains. The left two panels show 20 GM/WM means (ten subjects, two tissues per subject), and the right two panels show 90 AAL means (ten subjects, nine bilateral regions per subject). Scatter plots include the identity line; Bland–Altman plots (Bland & Altman, 1986) show the mean prediction–reference difference and 95% limits of agreement. Regions use the common reference/prediction reliability mask defined above. The scatter-panel R^2 denotes squared Pearson correlation r^2, whereas Table 1 uses coefficient-of-determination map agreement. Tissue and regional observations are repeated measurements within ten subjects.

On the common reference/prediction reliability mask, GM/WM means have an MAE of 1.63 ms and squared Pearson correlation $r^2 = 0.962$ across 20 subject–tissue measurements; AAL means have an MAE of 1.83 ms and $r^2 = 0.605$ across 90 subject–region measurements. These regional statistics use different aggregation and masking from the fixed-reference voxelwise R^2 in Table 1. Pooled correlations include between-tissue and between-region differences and do not isolate individual variation within a region. Predicted between-subject SDs are lower than their references (GM: 1.03 versus 2.13 ms; WM: 0.74 versus 1.80 ms; Appendix Table 10), consistent with attenuated subject variation.

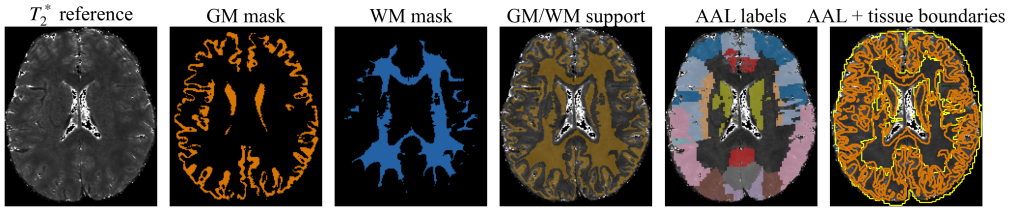


Figure 8: Segmentation and regional-support check for subject 06, slice 6, used in the regional agreement analysis. The overlays document the spatial support of the saved GM/WM and AAL summaries.

The scatter statistics in Figure 7 use squared Pearson correlation $r^2 = \mathrm{corr}(u, v)^2$. The map metric in Eq. (21) instead measures coefficient-of-determination agreement and penalizes offsets and scale errors. For regional differences $d = v - u$, descriptive Bland–Altman limits are $\bar{d} \pm 1.96 s_d$. Repeated regions within each participant are correlated, so these pooled limits are descriptive rather than participant-level confidence intervals.

Table 10: Regional reference and predicted means, subject-level absolute errors, and signed bias (prediction minus reference), in ms. Each row summarizes ten paired subject means; SDs use denominator nine.

Region	Reference	Prediction	MAE	Bias
GM	61.53 ± 2.13	60.02 ± 1.03	1.95 ± 0.98	-1.51 ± 1.63
WM	47.53 ± 1.80	46.68 ± 0.74	1.31 ± 0.85	-0.85 ± 1.36
Mid. cingulate	58.16 ± 1.53	57.59 ± 1.27	1.07 ± 0.99	-0.58 ± 1.37
Ant. cingulate	61.91 ± 2.05	60.99 ± 1.11	1.73 ± 1.09	-0.92 ± 1.89
Mid. frontal	60.03 ± 1.48	59.20 ± 1.43	1.59 ± 0.71	-0.83 ± 1.59
Sup. frontal	59.58 ± 2.22	58.37 ± 1.01	2.00 ± 0.80	-1.22 ± 1.85
Precentral	56.44 ± 2.02	54.96 ± 0.89	2.04 ± 1.54	-1.48 ± 2.14
Postcentral	56.73 ± 1.62	55.17 ± 0.77	1.97 ± 1.14	-1.56 ± 1.71
Supramarginal	58.14 ± 2.21	56.15 ± 0.76	2.11 ± 1.52	-1.99 ± 1.68
Precuneus	56.22 ± 2.02	54.79 ± 0.81	1.73 ± 1.34	-1.43 ± 1.69
Mid. occipital	55.73 ± 2.34	53.90 ± 1.56	2.24 ± 1.33	-1.82 ± 1.92

For the pooled descriptive agreement summaries, GM/WM has $N = 20$, MAE 1.63 ms, bias -1.18 ms, and limits $[-4.12, 1.76]$ ms; AAL has $N = 90$, MAE 1.83 ms, bias -1.32 ms, and limits $[-4.74, 2.11]$ ms. Here N counts repeated measurements from ten subjects.