RECONSTRUCT WHAT THE TASK NEEDS: PROGRESSIVE TASK-PRIOR INFORMATION BOTTLENECKS FOR ACCELERATED MRI

Anonymous authors
Paper under double-blind review

ABSTRACT

Magnetic resonance imaging (MRI) reconstruction is increasingly integrated with downstream tasks to preserve task-relevant structures. However, in multi-stage unrolled networks, stage-wise task feedback can entangle newly emerging task information with redundant and task-irrelevant content. We propose Purified Task Prior Guided MRI (PTP-MRI), which represents each stage update as a history-conditioned stochastic task increment. A variational lower bound encourages the retention of newly emerging task information, while a conditional variational upper bound limits task-irrelevant residual information. The resulting constrained information objective is optimized through Lagrangian duality with an adaptive dual variable. The purified task states are progressively accumulated and decoded into spatial priors that guide reconstruction according to task relevance. Experiments on SKM-TEA at $8\times$ acceleration and AHEAD at $16\times$ acceleration show that PTP-MRI improves downstream segmentation while maintaining reasonable reconstruction fidelity. Compared with the competing baseline for each metric, PTP-MRI achieves Dice/mIoU gains of 0.78/1.28 percentage points on SKM-TEA and 0.87/1.71 percentage points on AHEAD.

1 INTRODUCTION

Magnetic resonance imaging (MRI) reconstruction from undersampled k-space has been widely studied using deep cascaded and unrolled networks (Schlemper et al., 2017; Aggarwal et al., 2019; Sriram et al., 2020). Beyond image fidelity, recent task-driven methods further incorporate downstream objectives, such as segmentation or diagnosis, into reconstruction (Fan et al., 2018; Sun et al., 2019; Pramanik & Jacob, 2021; Karkalousos et al., 2024; Wu et al., 2024; He et al., 2025). Unrolled architectures are particularly suitable for such interaction, as task information can repeatedly influence intermediate reconstruction stages.

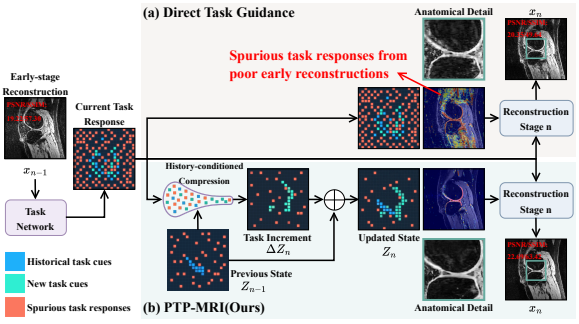

Figure 1: Motivation for PTP-MRI: purifying task guidance through history-conditioned information selection.

However, as illustrated in Fig. 1, progressive task feedback also creates an information-selection problem. Task responses at adjacent stages can be highly correlated, such that the current response may contain information already captured by previous stages together with only limited newly revealed task cues (Bansal et al., 2023; Peng et al., 2026). Moreover, responses derived from imperfect intermediate reconstructions may contain task-irrelevant variations induced by artifacts or noise (Wu et al., 2025; Gu et al., 2026; Deng et al., 2026). Propagating the complete response therefore mixes newly revealed task information with redundant and nuisance content (Zhang & Cheng, 2025; Kim et al., 2025). Effective task-driven reconstruction should instead determine *what is newly informative*, *what should be suppressed*, and *how the retained information should influence reconstruction*.

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.

To address this problem, we propose Purified Task Prior-Guided MRI Reconstruction (PTP-MRI). At each reconstruction stage, PTP-MRI extracts a stochastic task innovation conditioned on the accumulated task state, preserving newly revealed task information while constraining target-irrelevant residual information. The resulting constrained information objective is optimized through Lagrangian duality. The purified innovations are progressively accumulated and decoded into a spatial task prior, which modulates the learned image prior within the reconstruction objective. The main contributions are summarized as follows:

1) We formulate task-aware unrolled MRI reconstruction as a progressive information refinement problem, separating newly revealed task cues from redundant and task-irrelevant stage-wise responses.
2) We introduce a Task Prior Information Bottleneck (TPIB) that learns stochastic task innovations conditioned on the accumulated state. TPIB preserves task-relevant innovation while constraining target-irrelevant residual information, and solves the resulting constrained objective via Lagrangian dual optimization with an adaptive dual variable.
3) We propose task-prior-guided reconstruction, which decodes the progressive task state into a spatial prior and incorporates it into the reconstruction energy to modulate the learned image prior according to task relevance.

2 RELATED WORK

Deep Unrolled MRI Reconstruction. Deep MRI reconstruction has evolved from purely data-driven mapping toward model-based architectures that integrate learned priors with the MRI acquisition model. Cascaded networks (Schlemper et al., 2017) alternate learned image refinement with data-consistency operations, while unrolled methods such as MoDL (Aggarwal et al., 2019) and E2E-VarNet (Sriram et al., 2020) explicitly mimic iterative optimization across multiple reconstruction stages. Subsequent work has introduced Transformer-based modeling, generative priors, and neural operators to further improve reconstruction quality and robustness (Feng et al., 2023; Chung & Ye, 2022; Jatyani et al., 2025). These methods provide effective multi-stage reconstruction backbones, but are primarily optimized for image-domain fidelity rather than downstream task utility.

Task-Aware Reconstruction and Task Priors. Task-aware reconstruction incorporates downstream supervision or task-derived representations into the reconstruction process. Joint-learning methods couple reconstruction and downstream prediction through shared features, multi-task objectives, or stage-wise interactions (Sun et al., 2019; Pramanik & Jacob, 2021; Karkalousos et al., 2024; Wu et al., 2024), while other approaches feed segmentation, lesion, or anatomical priors back into reconstruction (Fan et al., 2018; He et al., 2025; Ju et al., 2026). These studies demonstrate the benefit of task guidance, but task responses are typically used as complete representations for feature fusion or feedback. Consequently, information inherited from earlier stages and variations arising from imperfect intermediate reconstructions may be propagated together with useful task cues.

Information Bottleneck for Representation Learning. The information bottleneck (IB) provides a principled framework for preserving task-relevant information while discarding irrelevant content (Tishby et al., 2000). Variational IB makes this objective tractable through stochastic latent variables and variational bounds (Alemi et al., 2017), while conditional variants constrain residual information given auxiliary context or targets (Fischer, 2020). Related ideas have also been explored in medical imaging for representation disentanglement and multimodal learning (Zuo et al., 2021; LIU et al., 2026). These formulations typically compress the current representation as a whole. In progressive unrolled reconstruction, however, the relevant question is not only what information should be retained, but what information is newly contributed beyond the accumulated state. This distinction motivates history-conditioned information selection across reconstruction stages.

3 METHOD

3.1 PROBLEM FORMULATION

Let $\mathbf{b}$ denote undersampled measurements acquired by a forward operator $\mathcal{A}$, and $\mathbf{x}$ the corresponding fully sampled image. Starting from $\hat{\mathbf{x}}_0$, an N-stage task-aware unrolled network with direct

task-response guidance takes the form

$$\mathbf{d}_n = \mathcal{D}_{\theta_n}(\hat{\mathbf{x}}_{n-1}), \qquad \mathbf{s}_n = \mathcal{T}_\omega(\mathbf{d}_n), \qquad \hat{\mathbf{x}}_n = \mathcal{U}(\hat{\mathbf{x}}_{n-1}, \mathbf{b}, \mathbf{d}_n, \mathbf{s}_n), \tag{1}$$

where $\mathcal{D}_{\theta_n}$ estimates an image prior, $\mathcal{T}_\omega$ extracts a task response, and $\mathcal{U}$ performs a physics-based reconstruction update. The full response $\mathbf{s}_n$ may contain previously captured information and task-irrelevant variations alongside new task cues, motivating history-conditioned information selection. Let S_n, Z_n, ΔZ_n, and Y denote the current task response, accumulated task state, stage-wise increment, and task target as random variables, with bold lowercase symbols denoting their realizations. For segmentation guidance, we use the foreground target $\mathbf{y} = \mathbf{1}[\mathbf{g} \neq 0]$, where $\mathbf{g}$ is the ground-truth segmentation mask. We model each task update as

$$\Delta\mathbf{z}_n \sim q_\phi(\Delta Z_n \mid \mathbf{s}_n, \mathbf{z}_{n-1}), \qquad \mathbf{z}_n = \mathbf{z}_{n-1} + \Delta\mathbf{z}_n, \qquad \mathbf{z}_0 = \mathbf{0}. \tag{2}$$

The increment should retain task information beyond the accumulated state while limiting task-irrelevant residual information. We formulate this objective as

$$\max_{\phi} \quad \sum_{n=1}^{N} I(Y; \Delta Z_n \mid Z_{n-1}) \qquad \text{s.t.} \quad \frac{1}{N}\sum_{n=1}^{N} I(S_n; \Delta Z_n \mid Y, Z_{n-1}) \leq C_R, \tag{3}$$

where $I(\cdot;\cdot \mid \cdot)$ denotes conditional mutual information and C_R is the residual information budget. The objective rewards task innovation beyond Z_{n-1}, while the constraint limits information retained from S_n beyond that explained by Y and Z_{n-1}.
To explicitly guide reconstruction, we decode the accumulated state into a spatial task prior:

$$\mathbf{t}_n = \sigma(\mathcal{C}_\xi(\mathbf{z}_n)), \qquad \hat{\mathbf{x}}_n = \mathcal{U}(\hat{\mathbf{x}}_{n-1}, \mathbf{b}, \mathbf{d}_n; \mathbf{t}_n), \tag{4}$$

where $\mathcal{C}_\xi$ is a learned decoder and σ is the sigmoid function. Solving this formulation requires tractable information objectives and an explicit mechanism for incorporating the spatial task prior into reconstruction.

3.2 Purified Task-Prior Guided MRI Reconstruction

To this end, we propose PTP-MRI, which combines a task prior information bottleneck (TPIB) with a task-prior-guided reconstruction update. TPIB learns progressive task states through variational approximations to the information objective and constraint in Eq. (3), while the reconstruction update incorporates the decoded spatial prior into a physics-based optimization problem.

3.2.1 Task Prior Information Bottleneck

As illustrated in Fig. 2, at stage n, TPIB takes the current task response $\mathbf{s}_n$ and the accumulated state $\mathbf{z}_{n-1}$ as inputs and produces a spatial task prior $\mathbf{t}_n$. To learn informative yet compact increments $\Delta\mathbf{z}_n$, we derive a variational lower bound on task innovation and a variational upper bound on task-irrelevant conditional information, as detailed below.

Variational task-innovation learning. The task-innovation term $I(Y; \Delta Z_n \mid Z_{n-1})$ measures task information contributed beyond the accumulated state. Since $Z_n = Z_{n-1} + \Delta Z_n$, we obtain

$$\begin{aligned} I(Y; \Delta Z_n \mid Z_{n-1}) &= H(Y \mid Z_{n-1}) - H(Y \mid Z_n, Z_{n-1}) \\ &\geq H(Y \mid Z_{n-1}) - H(Y \mid Z_n) \\ &\geq H(Y \mid Z_{n-1}) + \mathbb{E}\log p_\xi(Y \mid Z_n), \end{aligned} \tag{5}$$

where the equality follows from the equivalence of ΔZ_n and Z_n given Z_{n-1}. The inequalities follow from conditioning reducing entropy and the variational cross-entropy bound, respectively. Here, $p_\xi(Y \mid Z_n)$ is a learned task predictor. A detailed derivation is provided in the Appendix A.1. We stop gradients through the history state $\mathbf{z}_{n-1}$ when optimizing the stage-n innovation objective, treating its distribution as fixed for this local update. Thus, $H(Y \mid Z_{n-1})$ is treated as constant, and maximizing the variational lower bound reduces to minimizing the predictor's negative log-likelihood $-\mathbb{E}\log p_\xi(Y \mid Z_n)$.
Since Z_n is a latent random variable, its realization must be drawn from a learnable conditional distribution. Specifically, let $\tilde{\mathbf{z}}_{n-1} = \mathrm{sg}(\mathbf{z}_{n-1})$, where sg is the stop-gradient operator. We parameterize the conditional increment distribution as

$$q_\phi(\Delta Z_n \mid \mathbf{s}_n, \tilde{\mathbf{z}}_{n-1}) = \mathcal{N}\big(\boldsymbol{\mu}_n^{\Delta}, \mathrm{diag}\big((\boldsymbol{\sigma}_n^{\Delta})^2\big)\big), \tag{6}$$

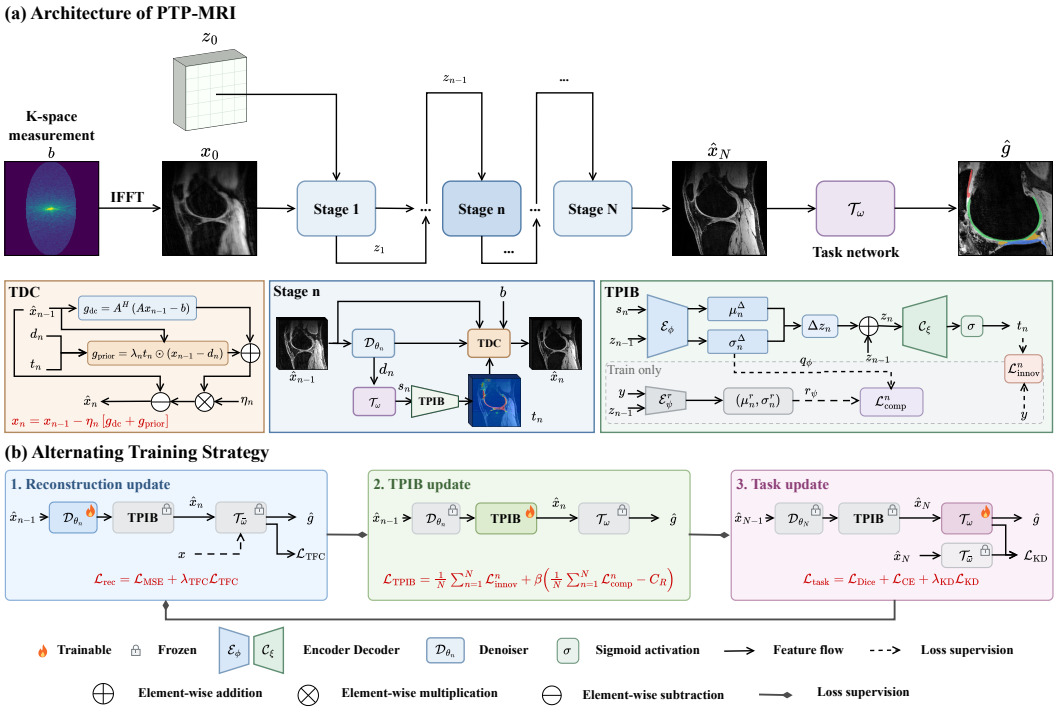

Figure 2: Overview of PTP-MRI. (a) At each reconstruction stage, TPIB extracts history-conditioned task increments and accumulates them into a spatial prior for task-guided data consistency. (b) Alternating training updates the reconstruction, TPIB, and task modules while keeping the other modules fixed.

with $(\boldsymbol{\mu}_n^{\Delta}, \boldsymbol{\sigma}_n^{\Delta}) = \mathcal{E}_\phi(\mathbf{s}_n, \tilde{\mathbf{z}}_{n-1})$, where $\mathcal{E}_\phi$ is the innovation encoder. We use reparameterization to sample the increment and update the state:

$$\Delta \mathbf{z}_n = \boldsymbol{\mu}_n^{\Delta} + \boldsymbol{\sigma}_n^{\Delta} \odot \boldsymbol{\epsilon}_n, \qquad \boldsymbol{\epsilon}_n \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \qquad \mathbf{z}_n = \tilde{\mathbf{z}}_{n-1} + \Delta \mathbf{z}_n. \tag{7}$$

Stop-gradient preserves the forward state accumulation while blocking gradient propagation through the history state.

The updated state is decoded into a spatial task prior:

$$\mathbf{t}_n = \sigma(\mathcal{C}_\xi(\mathbf{z}_n)), \tag{8}$$

where $\mathcal{C}_\xi$ is a learned decoder and σ denotes the sigmoid function. For the binary foreground target, we use a factorized Bernoulli predictor over M_G spatial locations:

$$p_\xi(\mathbf{y} \mid \mathbf{z}_n) = \prod_{j=1}^{M_G} t_{n,j}^{y_j} (1 - t_{n,j})^{1-y_j}. \tag{9}$$

The normalized variational negative log-likelihood is

$$\mathcal{L}_{\text{innov}}^n = -\frac{1}{M_G} \mathbb{E} \log p_\xi(Y \mid Z_n) = -\frac{1}{M_G} \mathbb{E} \sum_{j=1}^{M_G} [y_j \log t_{n,j} + (1 - y_j) \log(1 - t_{n,j})]. \tag{10}$$

The expectation is estimated using training samples and reparameterized latent increments. Under the fixed-history local update, minimizing $\mathcal{L}_{\text{innov}}^n$ maximizes the variational lower bound in Eq. (5).

Variational task-irrelevant compression. The innovation loss encourages ΔZ_n to retain task-relevant information, but does not prevent it from encoding additional variations in S_n that are irrelevant to Y. We therefore constrain $I(S_n; \Delta Z_n \mid Y, Z_{n-1})$, which measures the information about S_n retained in ΔZ_n after conditioning on the task target and the accumulated state. By definition,

$$I(S_n; \Delta Z_n \mid Y, Z_{n-1}) = \mathbb{E}\left[D_{\text{KL}}\left(q_\phi(\Delta Z_n \mid S_n, Y, Z_{n-1}) \,\|\, q_\phi(\Delta Z_n \mid Y, Z_{n-1})\right)\right]. \tag{11}$$

Since the innovation encoder generates ΔZ_n only from (S_n, Z_{n-1}), we have $\Delta Z_n \perp Y \mid (S_n, Z_{n-1})$, yielding

$$I(S_n; \Delta Z_n \mid Y, Z_{n-1}) = \mathbb{E}\left[D_{\mathrm{KL}}\left(q_\phi(\Delta Z_n \mid S_n, Z_{n-1}) \,\|\, q_\phi(\Delta Z_n \mid Y, Z_{n-1})\right)\right]. \tag{12}$$

The target-conditioned marginal

$$q_\phi(\Delta Z_n \mid Y, Z_{n-1}) = \int q_\phi(\Delta Z_n \mid S_n, Z_{n-1}) p(S_n \mid Y, Z_{n-1})\, dS_n \tag{13}$$

is generally intractable because $p(S_n \mid Y, Z_{n-1})$ is unknown. We therefore introduce a variational reference distribution $r_\psi(\Delta Z_n \mid Y, Z_{n-1})$, parameterized as a diagonal Gaussian:

$$r_\psi(\Delta Z_n \mid Y, Z_{n-1}) = \mathcal{N}\left(\boldsymbol{\mu}_n^r, \mathrm{diag}\left((\boldsymbol{\sigma}_n^r)^2\right)\right). \tag{14}$$

where $(\boldsymbol{\mu}_n^r, \boldsymbol{\sigma}_n^r) = \mathcal{E}_\psi^r(\mathbf{y}, \mathbf{z}_{n-1})$, with $\mathcal{E}_\psi^r(\cdot)$ denoting a reference encoder (Fig. 2). Replacing the intractable marginal with r_ψ gives

$$\begin{aligned}
&\mathbb{E}\left[D_{\mathrm{KL}}\left(q_\phi(\Delta Z_n \mid S_n, Z_{n-1}) \,\|\, r_\psi(\Delta Z_n \mid Y, Z_{n-1})\right)\right] \\
&= I(S_n; \Delta Z_n \mid Y, Z_{n-1}) + \mathbb{E}\left[D_{\mathrm{KL}}\left(q_\phi(\Delta Z_n \mid Y, Z_{n-1}) \,\|\, r_\psi(\Delta Z_n \mid Y, Z_{n-1})\right)\right] \\
&\geq I(S_n; \Delta Z_n \mid Y, Z_{n-1}),
\end{aligned} \tag{15}$$

where the inequality follows from the non-negativity of KL divergence. We therefore define the tractable compression loss as

$$\mathcal{L}_{\mathrm{comp}}^n = \mathbb{E}\left[D_{\mathrm{KL}}\left(q_\phi(\Delta Z_n \mid S_n, Z_{n-1}) \,\|\, r_\psi(\Delta Z_n \mid Y, Z_{n-1})\right)\right], \tag{16}$$

which upper-bounds the task-irrelevant information retained in ΔZ_n. Since both distributions are diagonal Gaussians, $\mathcal{L}_{\mathrm{comp}}^n$ admits the closed form

$$\mathcal{L}_{\mathrm{comp}}^n = \mathbb{E}\left[\frac{1}{2}\sum_{m=1}^{M}\left(\log\frac{(\sigma_{n,m}^r)^2}{(\sigma_{n,m}^\Delta)^2} + \frac{(\sigma_{n,m}^\Delta)^2 + (\mu_{n,m}^\Delta - \mu_{n,m}^r)^2}{(\sigma_{n,m}^r)^2} - 1\right)\right]. \tag{17}$$

where M is the total number of elements of $\Delta\mathbf{z}_n$. The expectation is approximated by averaging the closed-form KL divergence over the training mini-batch.

Using the variational bounds derived above, we replace the task-innovation objective with its negative log-likelihood surrogate and the conditional information constraint with its KL upper bound, yielding

$$\min_{\phi,\psi,\xi} \quad \frac{1}{N}\sum_{n=1}^{N}\mathcal{L}_{\mathrm{innov}}^n \qquad \text{s.t.} \quad \frac{1}{N}\sum_{n=1}^{N}\mathcal{L}_{\mathrm{comp}}^n \leq C_R. \tag{18}$$

Its Lagrangian formulation is

$$\min_{\phi,\psi,\xi}\max_{\beta\geq 0} \quad \mathcal{L}_{\mathrm{TPIB}} = \frac{1}{N}\sum_{n=1}^{N}\mathcal{L}_{\mathrm{innov}}^n + \beta\left(\frac{1}{N}\sum_{n=1}^{N}\mathcal{L}_{\mathrm{comp}}^n - C_R\right), \tag{19}$$

where β is the dual variable associated with the compression budget. We update the TPIB parameters by gradient descent under the local stop-gradient convention and update β by projected gradient ascent:

$$\beta \leftarrow \left[\beta + \eta_\beta\left(\frac{1}{N}\sum_{n=1}^{N}\widehat{\mathcal{L}}_{\mathrm{comp}}^n - C_R\right)\right]_+, \tag{20}$$

where η_β is the dual step size, $\widehat{\mathcal{L}}_{\mathrm{comp}}^n$ is the mini-batch estimate, and $[a]_+ = \max(a, 0)$.

3.2.2 Task-Prior-Guided Reconstruction

After getting spatial task prior $\mathbf{t}_n$ from TPIB, we instantiate the physics-based operator $\mathcal{U}$ in Eq. (4) as a task-driven data-consistency (TDC) update. Given the image prior $\mathbf{d}_n$ and spatial task prior $\mathbf{t}_n$, TDC balances measurement consistency with spatially weighted regularization toward $\mathbf{d}_n$. Holding both priors fixed during the update, we define the stage-wise energy as

$$\mathcal{E}_n(\mathbf{x}) = \frac{1}{2}\|\mathcal{A}\mathbf{x} - \mathbf{b}\|_2^2 + \frac{\lambda_n}{2}\left\|\sqrt{\mathbf{t}_n}\odot(\mathbf{x} - \mathbf{d}_n)\right\|_2^2, \tag{21}$$

where $\lambda_n > 0$ controls the regularization strength. The first term enforces consistency with the acquired measurements, while the second assigns stronger image-prior regularization to locations with higher predicted foreground probabilities. A gradient step on this energy yields the physics-based, task-guided reconstruction update:

$$\hat{\mathbf{x}}_n = \mathcal{U}(\hat{\mathbf{x}}_{n-1}, \mathbf{b}, \mathbf{d}_n; \mathbf{t}_n) = \hat{\mathbf{x}}_{n-1} - \eta_n \Big[\mathcal{A}^H(\mathcal{A}\hat{\mathbf{x}}_{n-1} - \mathbf{b}) + \lambda_n \mathbf{t}_n \odot (\hat{\mathbf{x}}_{n-1} - \mathbf{d}_n) \Big], \tag{22}$$

where $\mathcal{A}^H$ is the adjoint of $\mathcal{A}$ and $\eta_n > 0$ is the stage-wise step size. Thus, $\mathbf{d}_n$ specifies the preferred image content, while $\mathbf{t}_n$ spatially controls its influence on the physics-based reconstruction update.

3.3 Alternating Training and Inference

The reconstruction, TPIB, and task modules serve different objectives. We optimize them alternately within each training iteration.

Reconstruction update. The reconstruction module consists of a denoiser and a physics-based TDC update with learnable stage-specific step sizes and regularization weights. With TPIB and the task network fixed, we jointly optimize the denoiser parameters and these TDC parameters. With TPIB and the task network fixed, we optimize the Denoiser parameter θ_n using MSE loss and task feature consistency loss $\mathcal{L}_{\text{TFC}}$,

$$\mathcal{L}_{\text{rec}} = \mathcal{L}_{\text{MSE}} + \lambda_{\text{TFC}} \mathcal{L}_{\text{TFC}} = \frac{1}{N} \sum_{n=1}^{N} \|\hat{\mathbf{x}}_n - \mathbf{x}\|_2^2 + \frac{\lambda_{\text{TFC}}}{N} \sum_{n=1}^{N} \sum_{k \in \mathcal{K}} \left\| \mathcal{T}_\omega^k(\hat{\mathbf{x}}_n) - \mathcal{T}_\omega^k(\mathbf{x}) \right\|_1 . \tag{23}$$

Here, $\mathcal{T}_\omega^k(\cdot)$ denotes the feature representation at the k-th selected layer of the frozen task network.

TPIB update. We freeze $\{\theta_n\}_{n=1}^N$ and ω and update the innovation encoder, reference encoder, and task-prior decoder parameters (ϕ, ψ, ξ) by minimizing Eq. (19) with β fixed. We then update the dual variable β by projected gradient ascent. We apply stop-gradient to the previous task state so that each stage learns incremental information relative to a fixed history.

Task-network update. Finally, with the reconstruction network and TPIB fixed, we optimize the task network using the traditional Dice and cross entropy loss, as well as a knowledge distillation loss

$$\mathcal{L}_{\text{task}} = \mathcal{L}_{\text{Dice}} + \mathcal{L}_{\text{CE}} + \lambda_{\text{KD}} \mathcal{L}_{\text{KD}} \tag{24}$$

with $\mathcal{L}_{\text{KD}} = \sum_{k \in \mathcal{K}} \left\| \mathcal{T}_{\bar{\omega}}^k(\hat{\mathbf{x}}_N) - \mathcal{T}_\omega^k(\hat{\mathbf{x}}_N) \right\|_1$, which distills intermediate features from a fixed reference network $\mathcal{T}_{\bar{\omega}}$ trained on the fully sampled image $\mathbf{x}$, encouraging the task network to preserve task-relevant representations across reconstruction stages. Here, $\mathcal{K}$ denotes the selected feature levels and λ_{KD} balances the terms.

Inference. Sampling and the reference distribution r_ψ are required only during training. At inference, each stage uses the posterior mean to deterministically update $\bar{\mathbf{z}}_n = \bar{\mathbf{z}}_{n-1} + \boldsymbol{\mu}_n^\Delta$, constructs $\mathbf{t}_n$, and performs the task-prior-guided reconstruction update in Eq. (22). The final segmentation is

$$\hat{\mathbf{g}} = \arg\max_c p_{\omega^*}(c \mid \hat{\mathbf{x}}_N). \tag{25}$$

The complete training and inference procedures are summarized in Algorithm 1 of the Appendix A.3.

4 Experiments

Datasets. We evaluate on SKM-TEA (Desai et al., 2021) and AHEAD (Alkemade et al., 2020), covering knee and brain MRI. SKM-TEA contains 155 knee scans acquired with qDESS, with raw multi-coil k-space data and manual tissue annotations. We use both echoes and the provided 86/33/36 training/validation/test split. The segmentation targets are patellar cartilage, femoral cartilage, tibial cartilage, and meniscus, with medial and lateral labels merged for the latter two tissues. For AHEAD, we use a subset of 70 subjects from the 105-subject dataset, with 60 subjects for training and 10 for testing. Experiments use the inv1te1 contrast and three foreground classes: gray matter, white matter, and cerebrospinal fluid. Preprocessing and reference-label generation are detailed in Appendix A.2.

Baselines and experimental protocol. We compare against Medl (Qiao et al., 2023), MTrans (Feng et al., 2023), IDSLR (Pramanik & Jacob, 2021), A2MC-MRI-fix (He et al., 2025), and a neural operator (NO) model (Jatyani et al., 2025). A2MC-MRI-fix disables the original learnable sampler and uses the same fixed masks as the other methods. We use two-dimensional Gaussian variable-density random sampling with a fully sampled 24×24 central region; the same mask is used during training and testing for each sampling configuration. All methods share the same data splits, preprocessing, sampling masks, and evaluation metrics. The comparison methods, except IDSLR, follow a sequential protocol: the reconstruction network is trained first and then frozen while a common downstream segmentation network is trained. IDSLR follows its original end-to-end joint training protocol, in which the reconstruction and segmentation networks are optimized together. Their original loss formulations are retained for the corresponding training objectives. PTP-MRI uses the alternating training strategy described in Sec. 3.3. Implementation details for each baseline are provided in Appendix A.3.

Implementation and evaluation. PTP-MRI uses six reconstruction stages and is trained for 150 epochs using AdamW with cosine learning-rate annealing. Initial learning rates are 10^{-4} for the reconstruction and TPIB modules and 10^{-3} for the segmentation network. We evaluate reconstruction using PSNR and SSIM, and segmentation using Dice and mIoU. To characterize the information bottleneck, we additionally report information gain (IG), the reduction in segmentation cross-entropy after a task-state update, and residual gain (RG), the additional predictability of the task response provided by the current increment beyond the task label and historical state. Higher IG and lower RG indicate greater predictive improvement and reduced sensitivity to nuisance variations, respectively. Full training settings and metric definitions are provided in Appendices A.3 and A.4.

5 Results and Analysis

5.1 Comparison with Competing Methods

Fig. 3 compares reconstruction and downstream segmentation on SKM-TEA ($8\times$) and AHEAD ($16\times$). On SKM-TEA, IDSLR and MTrans exhibit blurring and structural distortions, while MedI, A2MC-fix, and NO produce visually plausible reconstructions but less accurate boundaries in thin structures. Our method better preserves cartilage and meniscus continuity, yielding contours closer to the ground truth and the highest Dice/mIoU ($81.05\%/68.82\%$) among the compared methods. On AHEAD, it similarly preserves cortical detail and tissue boundaries, achieving the highest Dice/mIoU ($88.01\%/79.87\%$). Notably, A2MC-fix achieves higher PSNR/SSIM on both datasets but lower segmentation accuracy, highlighting that global reconstruction fidelity does not necessarily imply better task performance. Additional qualitative and quantitative results on SKM-TEA ($16\times$) and AHEAD ($8\times$) are provided in the Appendix A.5.

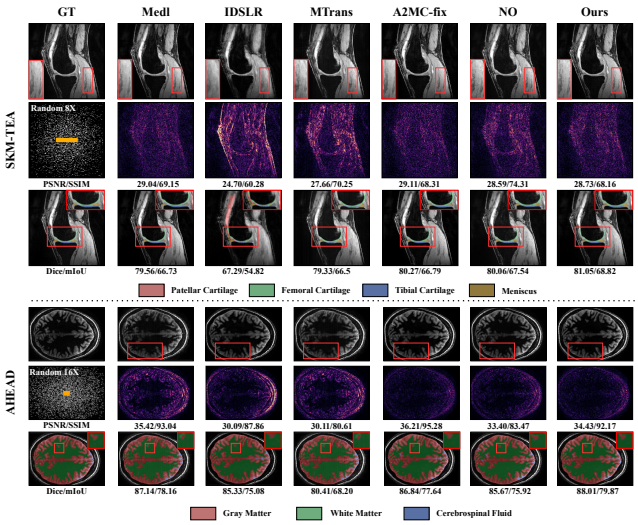

Figure 3: Reconstruction and segmentation results on two datasets obtained with different methods. Values below the error maps and segmentation results indicate PSNR/SSIM and Dice/mIoU, respectively.

5.2 Component Ablations and Further Analysis

Ablation of task innovation learning. We remove $\mathcal{L}_{\text{innov}}$, history conditioning (w/o hc), or state accumulation (w/o sa), keeping other components unchanged. The latter two variants use $q_\phi(\Delta Z_n \mid \mathbf{s}_n)$ and $Z_n = \Delta Z_n$, respectively. These ablations reduce Dice/mIoU by $0.36/0.71$, $0.68/1.12$, and $0.87/1.27$ percentage points despite comparable PSNR (Table 1). Fig. 4(a) shows

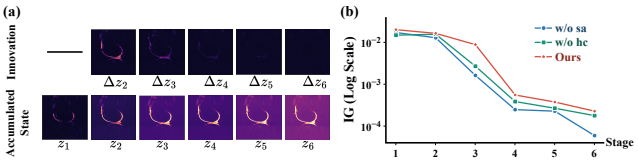

Figure 4: Task-innovation analysis. (a) stage-wise increments and accumulated states; (b) information gain across stages.

Table 1: Ablation of task-innovation learning. Dice, mIoU, and SSIM are reported as percentages; PSNR is reported in dB.

Variant	Dice ↑	mIoU ↑	PSNR ↑	SSIM ↑
w/o $\mathcal{L}_{\text{innov}}$	80.69	68.11	28.664	67.62
w/o hc	80.37	67.70	28.742	67.72
w/o sa	80.18	67.55	**28.759**	68.02
Full model	**81.05**	**68.82**	28.734	**68.16**

diminishing increment responses as the task state accumulates, consistent with decreasing IG across stagesin Fig. 4(b). The full model nevertheless maintains higher IG than w/o hc and w/o sa, indicating greater incremental predictive benefit beyond the available history. This supports the role of history conditioning and accumulation in extracting complementary task cues, helping explain the segmentation gains.

Ablation of task-irrelevant compression. We compare three variants under the same training protocol: removing $\mathcal{L}_{\text{comp}}$, replacing $r_\psi(\Delta Z_n \mid Y, Z_{n-1})$ with $\mathcal{N}(\mathbf{0}, \mathbf{I})$ in the KL term (Standard VIB), and using the learned conditional reference (full model). Table 2 shows that removing compression reduces Dice/mIoU by $0.52/0.89$ percentage points and PSNR by 0.77 dB. Standard VIB yields larger Dice/mIoU drops of $1.78/1.64$ points, despite higher PSNR than the variant without compression. The full model achieves the best performance across all four metrics, indicating that the choice of reference distribution matters beyond imposing compression alone. We further analyze how different compression strategies affect task-information retention and nuisance suppression through the IG–RG comparison in Fig. 5. Without compression, increments achieve high IG but also high RG, suggesting that useful task cues coexist with task-irrelevant variations. Standard VIB substantially lowers RG, but its lower IG suggests that compression against an unconditional reference also removes useful predictive information. In contrast, our model retains IG comparable to no compression while achieving RG close to Standard VIB. This combination indicates that the increments provide predictive information beyond the available history while remaining less sensitive to nuisance variations, helping explain the superior segmentation performance of the learned conditional reference.

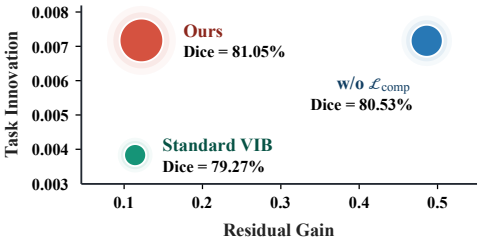

Figure 5: Task innovation (IG) versus residual gain (RG) under different compression strategies. Bubble size indicates Dice.

Table 2: Ablation of task-irrelevant compression at $8\times$ acceleration. Dice, mIoU, and SSIM are reported as percentages; PSNR is reported in dB.

Variant	Dice ↑	mIoU ↑	PSNR ↑	SSIM ↑
w/o $\mathcal{L}_{\text{comp}}$	80.53	67.93	27.96	67.49
Standard VIB	79.27	67.18	28.17	67.85
Full model	81.05	68.82	28.73	68.16

Ablation of task-guided spatial prior. We evaluate the spatial prior $\mathbf{t}_n$ by removing its guidance (NT), randomly permuting its pixels (PS), shuffling spatial blocks while preserving within-block structure (BS), or swapping priors between samples at corresponding stages (CS). Table 3 and Fig. 6 show that the spatial organization of task priors affects both reconstruction and segmentation. Ours achieves the best performance across all metrics, outperforming NT by 1.30 dB in global PSNR, 1.83 dB in ROI PSNR, and 1.38 percentage points in Dice. Disrupting the prior maps also degrades performance: CS retains more coherent spatial structure and performs better overall than BS and PS, but remains below the full model. These results suggest that effective guidance depends on the spatial correspondence between prior responses and target anatomy.

The visual results in Fig.6 support this interpretation. Our task spatial prior better preserves anatomical correspondence, whereas CS introduces spatial misalignment, BS produces blockwise discontinuities, and PS disrupts local coherence. These changes are accompanied by less faithful local

reconstruction and deviations in segmentation boundaries, particularly near the arrow-marked region. Our method better preserves thin tissue structures and produces segmentation contours closer to the ground truth, consistent with its superior ROI PSNR of 24.11 dB and ROI SSIM of 91.60%. Together, these findings support the importance of anatomically aligned spatial guidance for recovering task-relevant image details.

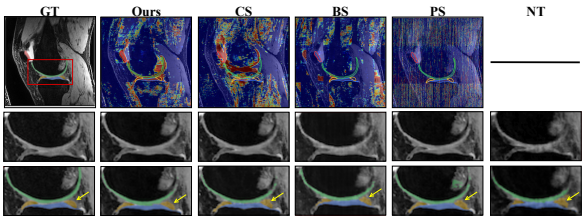

Figure 6: Different task prior maps and the corresponding reconstruction and segmentation results.

Table 3: Ablation of task-aware spatial guidance. Dice, mIoU, and SSIM are percentages; PSNR is in dB. ROI metrics are evaluated within target regions.

Variant	Dice ↑	mIoU ↑	PSNR ↑	SSIM ↑	PSNR_{ROI} ↑	SSIM_{ROI} ↑
PS	80.37	67.59	27.66	66.94	22.61	90.10
BS	80.42	67.63	27.70	66.84	22.79	90.27
CS	80.76	68.19	28.24	67.38	23.69	90.88
NT	79.67	65.92	27.43	66.07	22.28	89.87
Ours	**81.05**	**68.82**	**28.73**	**68.16**	**24.11**	**91.60**

Effect of the Information Budget. We vary the information budget as $C_R \in \{0.01, 0.02, 0.04\}$ and train each model to convergence, with all other settings fixed. For each budget, we report the converged dual variable β, the average compression upper bound $\bar{\mathcal{L}}_{\text{comp}} = \frac{1}{N}\sum_{n=1}^{N} \mathcal{L}_{\text{comp}}^{n}$, and the corresponding Dice and PSNR. As shown in Fig. 7, increasing C_R reduces the converged β from 0.2627 to 0.0870 and 0.0322, while the compression upper bound increases from 0.0102 to 0.0206 and 0.0437. The upper bound approximately tracks the prescribed budget, although it slightly exceeds the target, particularly at $C_R = 0.04$. These results indicate that the dual variable adapts the compression strength to the specified information budget. The resulting performance reveals a trade-off between reconstruction fidelity and downstream utility. Increasing C_R from 0.01 to 0.02 improves Dice from 80.69% to 81.05% and PSNR from 28.59 to 28.73 dB, suggesting that the tighter budget removes useful information. Further increasing C_R to 0.04 raises PSNR slightly to 28.78 dB but reduces Dice to 80.58%. This divergence is consistent with weaker compression retaining additional information that benefits image fidelity without improving tissue discrimination. Among the evaluated budgets, $C_R = 0.02$ achieves the highest Dice while maintaining PSNR close to the best observed value, supporting its use for task-aware reconstruction.

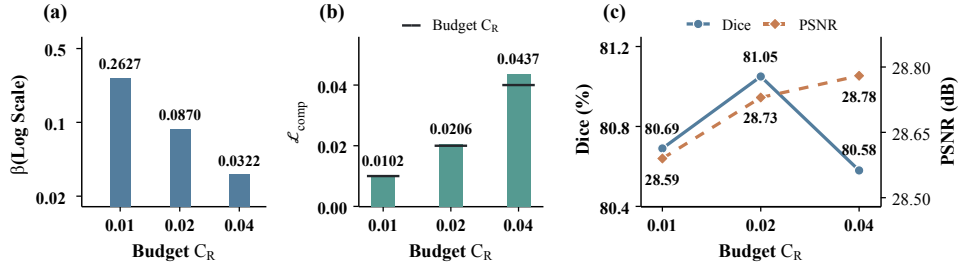

Figure 7: Effect of the information budget on (a) the dual variable, (b) the compression upper bound, and (c) reconstruction and segmentation performance.

6 Conclusion

We proposed PTP-MRI for task-aware unrolled MRI reconstruction. Instead of propagating the complete task response across stages, PTP-MRI explicitly models stage-wise task innovation conditioned on the accumulated task state. TPIB preserves newly revealed task information while suppressing target-irrelevant residuals under a constrained objective solved by Lagrangian dual optimization. The accumulated state is further decoded into a spatial task prior that modulates the learned image prior within the reconstruction update. Experimental results show that the proposed progressive information refinement improves downstream task performance while maintaining competitive image reconstruction quality. While the current framework relies on task supervision and is evaluated primarily for segmentation-guided reconstruction, extending progressive task information refinement to weaker supervision and more diverse downstream tasks is an important direction for future work.

AI USE STATEMENT

We used ChatGPT for non-scientific writing assistance, including grammar correction, sentence rephrasing, and improving clarity in the introduction and discussion sections. We also used it to help debug implementation details of the released code. No AI tool was used to generate novel ideas, experimental results, figures, or references. All AI-assisted outputs were carefully checked and edited by the authors, who retain full responsibility for the manuscript.

REFERENCES

Hemant K. Aggarwal, Merry P. Mani, and Mathews Jacob. Modl: Model-based deep learning architecture for inverse problems. *IEEE Transactions on Medical Imaging*, 38(2):394–405, 2019. doi: 10.1109/TMI.2018.2865356.

Alexander A. Alemi, Ian Fischer, Joshua V. Dillon, and Kevin Murphy. Deep variational information bottleneck. In *International Conference on Learning Representations*, 2017. URL https://openreview.net/forum?id=HyxQzBceg.

Anneke Alkemade, Martijn J. Mulder, Josephine M. Groot, Bethany R. Isaacs, Nikita van Berendonk, Nicky Lute, Scott J. S. Isherwood, Pierre-Louis Bazin, and Birte U. Forstmann. The amsterdam ultra-high-field adult lifespan database (AHEAD): A freely available multimodal 7 tesla submillimeter magnetic resonance imaging database. *NeuroImage*, 221:117200, November 2020. doi: 10.1016/j.neuroimage.2020.117200.

Dishank Bansal, Ricky TQ Chen, Mustafa Mukadam, and Brandon Amos. Taskmet: Task-driven metric learning for model learning. *Advances in Neural Information Processing Systems*, 36: 46505–46519, 2023.

Hyungjin Chung and Jong Chul Ye. Score-based diffusion models for accelerated MRI. *Medical Image Analysis*, 80:102479, August 2022. doi: 10.1016/j.media.2022.102479.

Zeyu Deng, Lihui Wang, Xi Tao, Qijian Chen, Ying Cao, Yingfeng Ou, et al. Uncertainty-propelled physics-mae fusion for self-supervised diffusion-weighted image denoising. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 3578–3586, 2026.

Arjun D. Desai, Andrew M. Schmidt, Elka B. Rubin, Christopher M. Sandino, Marianne S. Black, Valentina Mazzoli, Kathryn J. Stevens, Robert D. Boutin, Christopher Ré, Garry E. Gold, Brian A. Hargreaves, and Akshay S. Chaudhari. SKM-TEA: A dataset for accelerated MRI reconstruction with dense image labels for quantitative clinical evaluation. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1. Curran Associates, Inc., 2021. URL https://datasets-benchmarks-proceedings.neurips.cc/paper/2021/hash/03c6b06952c750899bb03d998e631860-Abstract-round2.html.

Zhiwen Fan, Liyan Sun, Xinghao Ding, Yue Huang, Congbo Cai, and John Paisley. A segmentation-aware deep fusion network for compressed sensing MRI. In *Computer Vision – ECCV 2018*, volume 11210 of *Lecture Notes in Computer Science*, pp. 55–70. Springer, 2018. doi: 10.1007/978-3-030-01231-1_4.

Chun-Mei Feng, Yunlu Yan, Geng Chen, Yong Xu, Ying Hu, Ling Shao, and Huazhu Fu. Multimodal transformer for accelerated mr imaging. *IEEE Transactions on Medical Imaging*, 42(10): 2804–2816, 2023. doi: 10.1109/TMI.2022.3180228.

Ian Fischer. The conditional entropy bottleneck. *Entropy*, 22(9):999, September 2020. doi: 10.3390/e22090999.

Xianming Gu, Lihui Wang, Ying Cao, Zeyu Deng, Yingfeng Ou, Guodong Hu, and Yi Chen. Cddpe: Dual-prompt expert network based on convolutional dictionary feature decoupling for multi-contrast mri super-resolution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 4329–4338, 2026.

Yuzhu He, Chunfeng Lian, Ruyi Xiao, Fangmao Ju, Chao Zou, Zongben Xu, and Jianhua Ma. Anatomy-aware deep unrolling for task-oriented acceleration of multi-contrast MRI. *IEEE Transactions on Medical Imaging*, 44(9):3832–3844, September 2025. doi: 10.1109/TMI.2025.3568157.

Armeet Singh Jatyani, Jiayun Wang, Aditi Chandrashekar, Zihui Wu, Miguel Liu-Schiaffini, Bahareh Tolooshams, and Anima Anandkumar. A unified model for compressed sensing mri across undersampling patterns. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 26004–26013, June 2025.

Fangmao Ju, Yuzhu He, Fan Wang, Xianjun Li, Chen Niu, Chunfeng Lian, and Jianhua Ma. Model-unrolled fast MRI with weakly supervised lesion enhancement. *Medical Image Analysis*, 107: 103806, January 2026. doi: 10.1016/j.media.2025.103806.

Dimitrios Karkalousos, Ivana Isgum, Henk Marquering, and Matthan W. A. Caan. MultiTask learning for accelerated-MRI reconstruction and segmentation of brain lesions in multiple sclerosis. In *Medical Imaging with Deep Learning (MIDL)*, volume 227 of *Proceedings of Machine Learning Research*, pp. 991–1005. PMLR, July 2024. URL https://proceedings.mlr.press/v227/karkalousos24a.html.

Jaeha Kim, Junghun Oh, and Kyoung Mu Lee. Exploiting diffusion prior for task-driven image restoration. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 10151–10161. IEEE, 2025.

TIANYI LIU, Xi Yang, Wei Wang, Anh Nguyen, Haochuan Jiang, and Kaizhu Huang. Sequential information bottleneck fusion: Towards robust and generalizable multi-modal brain tumor segmentation. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=tmV2sOZ8TV.

Xinyu Peng, Ziyang Zheng, Wenrui Dai, Duoduo Xue, Shaohui Li, Chenglin Li, Junni Zou, and Hongkai Xiong. Information-theoretic optimization for task-adapted compressed sensing magnetic resonance imaging. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 48 (8):9947–9965, 2026. doi: 10.1109/TPAMI.2026.3683201.

Aniket Pramanik and Mathews Jacob. Reconstruction and segmentation of parallel mr data using image domain deep-slr. In *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pp. 1095–1098, 2021. doi: 10.1109/ISBI48211.2021.9434056.

Xiaoyu Qiao, Yuping Huang, and Weisheng Li. MEDL-Net: A model-based neural network for MRI reconstruction with enhanced deep learned regularizers. *Magnetic Resonance in Medicine*, 89(5):2062–2075, May 2023. doi: 10.1002/mrm.29575.

Jo Schlemper, Jose Caballero, Joseph Hajnal, Anthony Price, and Daniel Rueckert. A deep cascade of convolutional neural networks for dynamic mr image reconstruction. *IEEE Transactions on Medical Imaging*, 37:491–503, 04 2017. doi: 10.1109/TMI.2017.2760978.

Anuroop Sriram, Jure Zbontar, Tullie Murrell, Aaron Defazio, C. Lawrence Zitnick, Nafissa Yakubova, Florian Knoll, and Patricia Johnson. End-to-end variational networks for accelerated MRI reconstruction. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*, volume 12262 of *Lecture Notes in Computer Science*, pp. 64–73. Springer, 2020. doi: 10.1007/978-3-030-59713-9_7.

Liyan Sun, Zhiwen Fan, Xinghao Ding, Yue Huang, and John Paisley. Joint CS-MRI reconstruction and segmentation with a unified deep network. In *Information Processing in Medical Imaging (IPMI)*, volume 11492 of *Lecture Notes in Computer Science*, pp. 492–504. Springer, 2019. doi: 10.1007/978-3-030-20351-1_38.

Naftali Tishby, Fernando C Pereira, and William Bialek. The information bottleneck method. *arXiv preprint physics/0004057*, 2000.

Linhai Wu, Lihui Wang, Zeyu Deng, Yuemin Zhu, and Hongjiang Wei. Replace2self: Self-supervised denoising based on voxel replacing and image mixing for diffusion mri. *IEEE Transactions on Medical Imaging*, 44(7):2878–2891, 2025. doi: 10.1109/TMI.2025.3552611.

Zihui Wu, Tianwei Yin, Yu Sun, Robert Frost, Andre van der Kouwe, Adrian V. Dalca, and Katherine L. Bouman. Learning task-specific strategies for accelerated mri. *IEEE Transactions on Computational Imaging*, 10:1040–1054, 2024. doi: 10.1109/TCI.2024.3410521.

Dong Zhang and Kwang-Ting Cheng. Generalized task-driven medical image quality enhancement with gradient promotion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47 (4):2785–2798, 2025. doi: 10.1109/TPAMI.2025.3525671.

Lianrui Zuo, Blake E. Dewey, Yihao Liu, Yufan He, Scott D. Newsome, Ellen M. Mowry, Susan M. Resnick, Jerry L. Prince, and Aaron Carass. Unsupervised MR harmonization by learning disentangled representations using information bottleneck theory. *NeuroImage*, 243:118569, November 2021. doi: 10.1016/j.neuroimage.2021.118569.

A Appendix

A.1 Derivation of the Task-Innovation Variational Bound

Let Y be the task variable, and suppose that the accumulated state satisfies $Z_n = Z_{n-1} + \Delta Z_n$. All expectations below are taken under the joint distribution induced by the data and the model. We assume that the relevant conditional entropies and cross-entropies are finite.

Conditional mutual information. By definition,

$$I(Y; \Delta Z_n \mid Z_{n-1}) = H(Y \mid Z_{n-1}) - H(Y \mid \Delta Z_n, Z_{n-1}). \tag{A.1}$$

Given Z_{n-1}, the mapping between ΔZ_n and $Z_n = Z_{n-1} + \Delta Z_n$ is invertible. Hence

$$I(Y; \Delta Z_n \mid Z_{n-1}) = H(Y \mid Z_{n-1}) - H(Y \mid Z_n, Z_{n-1}). \tag{A.2}$$

Since

$$H(Y \mid Z_n) - H(Y \mid Z_n, Z_{n-1}) = I(Y; Z_{n-1} \mid Z_n) \geq 0, \tag{A.3}$$

we obtain

$$I(Y; \Delta Z_n \mid Z_{n-1}) \geq H(Y \mid Z_{n-1}) - H(Y \mid Z_n). \tag{A.4}$$

Variational bound. Let $p_\xi(Y \mid Z_n)$ be a task predictor whose support contains that of the true conditional distribution $p(Y \mid Z_n)$. The conditional cross-entropy identity gives

$$\begin{aligned} -\mathbb{E} \log p_\xi(Y \mid Z_n) = H(Y \mid Z_n) \\ + \mathbb{E}_{Z_n} \mathrm{KL}(p(\cdot \mid Z_n) \,\|\, p_\xi(\cdot \mid Z_n)) \,. \end{aligned} \tag{A.5}$$

The KL divergence is non-negative; substituting the resulting upper bound on $H(Y \mid Z_n)$ into equation A.4 yields

$$I(Y; \Delta Z_n \mid Z_{n-1}) \geq H(Y \mid Z_{n-1}) + \mathbb{E} \log p_\xi(Y \mid Z_n). \tag{A.6}$$

Tightness of the bound. Combining equation A.2, equation A.3, and equation A.5 gives the exact identity

$$\begin{aligned} I(Y; \Delta Z_n \mid Z_{n-1}) = H(Y \mid Z_{n-1}) + \mathbb{E} \log p_\xi(Y \mid Z_n) \\ + \mathbb{E}_{Z_n} \mathrm{KL}(p(\cdot \mid Z_n) \,\|\, p_\xi(\cdot \mid Z_n)) \\ + I(Y; Z_{n-1} \mid Z_n). \end{aligned} \tag{A.7}$$

Thus, the bound is tight precisely when the predictor matches $p(Y \mid Z_n)$ almost surely and $I(Y; Z_{n-1} \mid Z_n) = 0$.

Implications for optimization. If the distribution of (Y, Z_{n-1}) is fixed with respect to the parameters being optimized, then $H(Y \mid Z_{n-1})$ is constant. Maximizing the right-hand side of equation A.6 is then equivalent to minimizing

$$\mathcal{L}_n = \mathbb{E}[-\log p_\xi(Y \mid Z_n)]. \tag{A.8}$$

Stopping gradients through Z_{n-1} gives the same local gradient with respect to parameters of the current update. It does not, however, make $H(Y \mid Z_{n-1})$ constant throughout joint training: earlier states can change as shared or preceding-stage parameters are updated. Consequently, cross-entropy minimization alone does not generally maximize the full right-hand side of equation A.6 during joint optimization.

An auxiliary predictor $q_\psi(Y \mid Z_{n-1})$ may be used to estimate $H(Y \mid Z_{n-1})$. Its cross-entropy satisfies

$$\begin{aligned} \mathbb{E}[-\log q_\psi(Y \mid Z_{n-1})] = H(Y \mid Z_{n-1}) \\ + \mathbb{E}_{Z_{n-1}} \mathrm{KL}(p(\cdot \mid Z_{n-1}) \,\|\, q_\psi(\cdot \mid Z_{n-1})) \,. \end{aligned} \tag{A.9}$$

Because this cross-entropy is an *upper* bound on the conditional entropy, substituting it into equation A.6 does not, in general, preserve a guaranteed lower bound on $I(Y; \Delta Z_n \mid Z_{n-1})$.

A.2 Dataset

SKM-TEA. We use the provided subject-level split of 86 training, 33 validation, and 36 test scans. Both qDESS echoes are included in the experiments. The medial and lateral tibial cartilage labels are merged into a single tibial cartilage class, and the medial and lateral meniscus labels are similarly merged. Together with patellar cartilage, femoral cartilage, and background, these form a five-class segmentation task.

AHEAD. We use the preprocessed inv1te1 contrast from 70 subjects, comprising 60 training and 10 test subjects. Multi-coil k-space data and sensitivity maps are compressed from 32 physical coils to 8 virtual coils using PCA. The central 200 slices of each volume are retained, yielding 12,000 training and 2,000 test slices of size 288×288. Reference tissue labels are generated using SPM12 and supplemented with atlas-defined subcortical regions transformed from MNI space to native space. The labels are consolidated into gray matter, white matter, and cerebrospinal fluid, with mapped subcortical gray-matter structures included in the gray-matter class. Background is included as the fourth class.

A.3 Implementation and Training Details

Component architectures and optimization. The reconstruction network comprises six stages. Figure A.1 presents the detailed architectures of selected components in PTP-MRI: the innovation encoder, reference (compression) encoder, task-prior decoder, and residual U-Net denoiser.

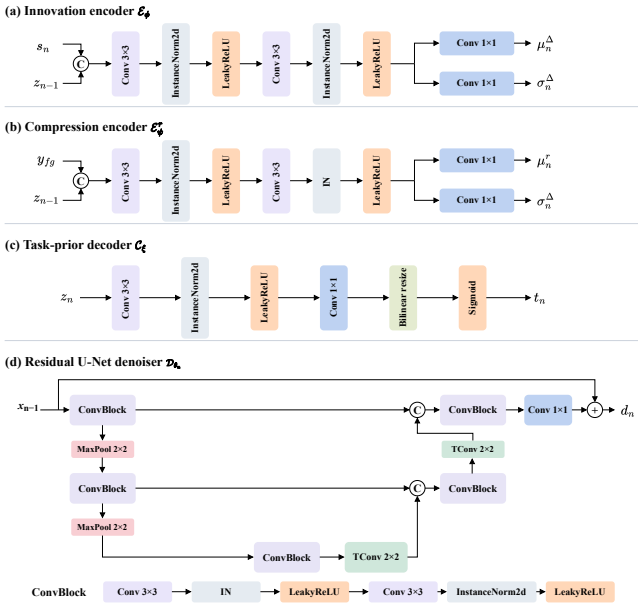

Figure A.1: Detailed architectures of selected components in PTP-MRI: (a) the innovation encoder, (b) the reference (compression) encoder, (c) the task-prior decoder, and (d) the residual U-Net denoiser.

We train for 150 epochs with a batch size of 16, using AdamW with weight decay 10^{-5} and cosine learning-rate annealing. Initial learning rates are 10^{-4} for the reconstruction and TPIB modules and 10^{-3} for the task network. Each experiment is repeated with three different random seeds (42, 100, and 88), and the mean performance across the three runs is reported.

Training Algorithm. Algorithm 1 summarizes the alternating optimization procedure and the corresponding deterministic inference scheme. During training, the reconstruction, TPIB, and task parameters are updated in separate steps, while the reference parameters $\bar{\omega}$ remain fixed. At inference time, sampled innovations are replaced by their conditional means.

Training objectives. The reconstruction objective is

$$\mathcal{L}_{\text{rec}} = \mathcal{L}_{\text{MSE}} + 0.01\,\mathcal{L}_{\text{TFC}}. \tag{A.10}$$

The task objective is

$$\mathcal{L}_{\text{task}} = \mathcal{L}_{\text{Dice}} + \mathcal{L}_{\text{CE}} + 0.01\,\mathcal{L}_{\text{KD}}. \tag{A.11}$$

The TPIB objective is

$$\mathcal{L}_{\text{TPIB}} = \frac{1}{N}\sum_{n=1}^{N} \mathcal{L}^{n}_{\text{innov}} + \beta \left(\frac{1}{N}\sum_{n=1}^{N} \mathcal{L}^{n}_{\text{comp}} - C_R \right). \tag{A.12}$$

The dual variable β is initialized to 1 and updated with a step size of 10^{-3}. Each alternating update optimizes the active module while keeping the other modules fixed.

Algorithm 1 Alternating Training and Deterministic Inference

Require: Training tuples $(\mathbf{b}, \mathbf{x}, \mathbf{g})$; number of stages N; trainable parameters $\{\theta_n\}_{n=1}^{N}$, (ϕ, ψ, ξ), and ω; fixed reference parameters $\bar{\omega}$

Training
for each training iteration **do**
  Sample a mini-batch of training tuples $(\mathbf{b}, \mathbf{x}, \mathbf{g})$
  Reconstruction update:
  Freeze (ϕ, ψ, ξ) and ω; enable updates to $\{\theta_n\}_{n=1}^{N}$
  Run the N-stage reconstruction and update $\{\theta_n\}_{n=1}^{N}$ by minimizing $\mathcal{L}_{\mathrm{rec}}$
  TPIB update:
  Freeze $\{\theta_n\}_{n=1}^{N}$ and ω; enable updates to (ϕ, ψ, ξ)
  Initialize $\mathbf{x}_0$ from $\mathbf{b}$ and set $\mathbf{z}_0 = \mathbf{0}$
  for $n = 1$ to N **do**
    $\mathbf{d}_n = \mathcal{D}_{\theta_n}(\mathbf{x}_{n-1})$, $\mathbf{s}_n = \mathcal{T}_{\omega}(\mathbf{d}_n)$
    Sample $\Delta\mathbf{z}_n \sim q_\phi(\Delta Z_n \mid \mathbf{s}_n, \mathbf{z}_{n-1})$
    $\mathbf{z}_n = \mathrm{sg}(\mathbf{z}_{n-1}) + \Delta\mathbf{z}_n$
    $\mathbf{t}_n = \sigma(\mathcal{C}_\xi(\mathbf{z}_n))$
    $\mathbf{x}_n = \mathrm{TDC}(\mathbf{x}_{n-1}, \mathbf{d}_n, \mathbf{t}_n, \mathbf{b})$
  end for
  Update (ϕ, ψ, ξ) by minimizing $\mathcal{L}_{\mathrm{TPIB}}$
  Task update:
  Freeze $\{\theta_n\}_{n=1}^{N}$ and (ϕ, ψ, ξ); enable updates to ω
  Update ω by minimizing $\mathcal{L}_{\mathrm{task}}$, with $\bar{\omega}$ held fixed
end for
Deterministic inference
Input: Measurement $\mathbf{b}$
Initialize $\mathbf{x}_0$ from $\mathbf{b}$ and set $\mathbf{z}_0 = \mathbf{0}$
for $n = 1$ to N **do**
  $\mathbf{d}_n = \mathcal{D}_{\theta_n}(\mathbf{x}_{n-1})$, $\mathbf{s}_n = \mathcal{T}_{\omega}(\mathbf{d}_n)$
  $\boldsymbol{\mu}_n^{\Delta} = \mathcal{E}_\phi(\mathbf{s}_n, \mathbf{z}_{n-1})$
  $\mathbf{z}_n = \mathbf{z}_{n-1} + \boldsymbol{\mu}_n^{\Delta}$
  $\mathbf{t}_n = \sigma(\mathcal{C}_\xi(\mathbf{z}_n))$
  $\mathbf{x}_n = \mathrm{TDC}(\mathbf{x}_{n-1}, \mathbf{d}_n, \mathbf{t}_n, \mathbf{b})$
end for
$\hat{\mathbf{G}} = \arg\max_c p_\omega(c \mid \mathbf{x}_N)$
Output: Reconstructed image $\mathbf{x}_N$ and task prediction $\hat{\mathbf{G}}$

Baseline implementation. Comparison methods are trained sequentially. For all comparison methods except IDSLR, we first train the reconstruction model, freeze its parameters, and then train a common segmentation network on its reconstructed images. IDSLR retains its original end-to-end joint training protocol, with the reconstruction and segmentation networks optimized together. The initial learning rates are 10^{-4} for reconstruction and 10^{-3} for segmentation. We retain the loss formulations of the original methods for the implemented training objectives. For A2MC-MRI-fix, the learnable sampling component is disabled and replaced with the shared fixed mask.

Computing environment. All experiments are implemented in PyTorch and run on a workstation equipped with a single NVIDIA RTX PRO 6000 Blackwell GPU and an Intel Xeon Gold 6530 CPU.

A.4 Evaluation metrics

We evaluate the information retained in the progressive task states using two complementary probe-based metrics: information gain (IG) and residual gain (RG). At stage n, let $\mathbf{s}_n$ denote the task response, $\mathbf{z}_{n-1}$ the historical task state, $\Delta\mathbf{z}_n$ the current increment, and $\mathbf{y}$ the task label. IG evaluates the additional task-predictive information provided by $\Delta\mathbf{z}_n$, whereas RG evaluates its contribution to predicting $\mathbf{s}_n$ after conditioning on $\mathbf{y}$ and $\mathbf{z}_{n-1}$.

After training, we freeze the evaluated model and extract $(\mathbf{s}_n, \mathbf{z}_{n-1}, \Delta\mathbf{z}_n, \mathbf{y})$ at each stage. For each metric, we train a pair of probes with comparable capacity and identical optimization settings using the training split. The probes are then fixed and evaluated on the test split, without updating the evaluated model. All losses below denote averages over the test samples, with cross-entropy averaged over spatial locations and MSE averaged over response elements.

Information gain (IG). We train two task-prediction probes: $P_A^{(n)}$ predicts $\mathbf{y}$ from the historical state $\mathbf{z}_{n-1}$, while $P_B^{(n)}$ additionally receives the current increment $\Delta\mathbf{z}_n$. Their evaluation losses are

$$L_{\text{hist}}^{(n)} = \mathbb{E}_{\text{test}}\left[\text{CE}\left(\mathbf{y}, P_A^{(n)}(\mathbf{z}_{n-1})\right)\right], \tag{A.13}$$

$$L_{\text{hist+inc}}^{(n)} = \mathbb{E}_{\text{test}}\left[\text{CE}\left(\mathbf{y}, P_B^{(n)}(\mathbf{z}_{n-1}, \Delta\mathbf{z}_n)\right)\right]. \tag{A.14}$$

The stage-wise information gain is

$$\text{IG}_n = L_{\text{hist}}^{(n)} - L_{\text{hist+inc}}^{(n)}. \tag{A.15}$$

Higher IG indicates that the increment improves task prediction beyond the information available in the historical state.

Residual gain (RG). We train a separate pair of response-prediction probes. $Q_A^{(n)}$ predicts $\mathbf{s}_n$ from $(\mathbf{y}, \mathbf{z}_{n-1})$, while $Q_B^{(n)}$ additionally receives $\Delta\mathbf{z}_n$. Using mean squared error, we compute

$$L_{\text{base}}^{(n)} = \mathbb{E}_{\text{test}}\left[\text{MSE}\left(\mathbf{s}_n, Q_A^{(n)}(\mathbf{y}, \mathbf{z}_{n-1})\right)\right], \tag{A.16}$$

$$L_{\text{inc}}^{(n)} = \mathbb{E}_{\text{test}}\left[\text{MSE}\left(\mathbf{s}_n, Q_B^{(n)}(\mathbf{y}, \mathbf{z}_{n-1}, \Delta\mathbf{z}_n)\right)\right]. \tag{A.17}$$

The stage-wise residual gain is

$$\text{RG}_n = L_{\text{base}}^{(n)} - L_{\text{inc}}^{(n)}. \tag{A.18}$$

A larger RG indicates that the increment makes more response information recoverable beyond the task label and history. Lower RG, when accompanied by high IG, supports suppression of conditional residual information while retaining useful task-predictive information.

A.5 Additional Results

Quantitative results. Tables A.1 and A.2 provide additional results on AHEAD at $8\times$ acceleration and SKM-TEA at $16\times$ acceleration, respectively. PTP-MRI achieves the highest Dice and mIoU among the evaluated reconstruction methods in both settings. On AHEAD, it improves Dice and mIoU by 1.02 and 0.61 percentage points over A2MC-fix, respectively. On SKM-TEA, the corresponding improvements over Medl are 0.59 and 0.70 percentage points. These segmentation gains are accompanied by lower PSNR and SSIM than the respective best-performing baselines, highlighting a trade-off between downstream task performance and global reconstruction fidelity.

Table A.1: Additional quantitative results on AHEAD at $8\times$ acceleration. Dice, mIoU, and SSIM are reported as percentages; PSNR is reported in dB. Bold indicates the best result among reconstruction methods. Oracle and No Reconstruction are reference settings. Our results are highlighted in light gray.

	Segmentation		Reconstruction	
Method	Dice ↑	mIoU ↑	PSNR ↑	SSIM ↑
Oracle	90.38	83.52	–	–
No Reconstruction	84.35	73.90	–	–
Medl	88.81	80.84	**38.11**	95.86
MTrans	84.42	73.82	31.53	82.90
IDSLR	88.27	79.58	32.88	92.06
A2MC-fix	88.82	80.87	38.04	**96.01**
NO	88.32	80.05	35.95	80.67
PTP-MRI (Ours)	**89.84**	**81.48**	36.48	95.33

Table A.2: Additional quantitative results on SKM-TEA at $16\times$ acceleration. Dice, mIoU, and SSIM are reported as percentages; PSNR is reported in dB. Bold indicates the best result among reconstruction methods. Oracle and No Reconstruction are reference settings. Our results are highlighted in light gray.

	Segmentation		Reconstruction	
Method	Dice ↑	mIoU ↑	PSNR ↑	SSIM ↑
Oracle	80.91	68.61	–	–
No Reconstruction	75.19	60.90	–	–
Medl	78.94	65.87	27.29	61.02
MTrans	76.40	62.62	26.41	65.54
IDSLR	48.81	37.25	23.82	52.97
A2MC-fix	78.08	64.86	**27.33**	59.12
NO	77.20	63.64	26.80	**69.31**
PTP-MRI (Ours)	**79.53**	**66.57**	27.00	59.66

Qualitative comparisons. Figure A.2 presents reconstructed images, error maps, and segmentation overlays for the two additional settings. In the SKM-TEA example, IDSLR exhibits pronounced blurring, while MTrans shows residual artifacts; PTP-MRI provides a sharper reconstruction with clearer cartilage delineation. For AHEAD, the segmentation overlays reveal differences in tissue delineation even among methods with visually similar reconstructions. The error maps and overlays complement the quantitative comparisons by showing where reconstruction and segmentation differences occur. The scores annotated in the figure report dataset-level performance.

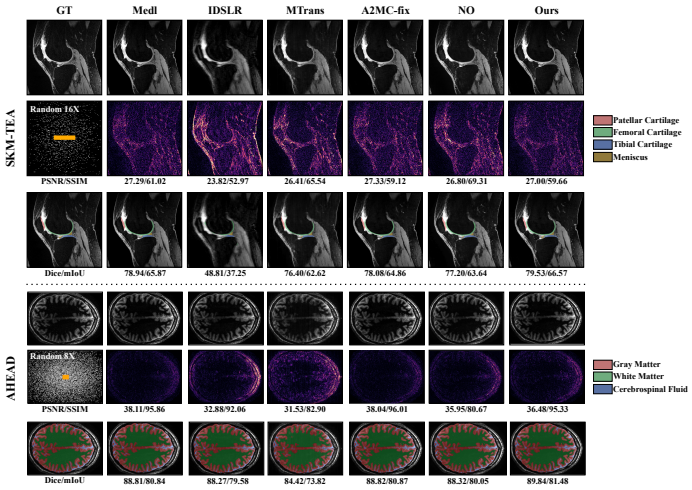

Figure A.2: Additional qualitative comparisons on SKM-TEA at $16\times$ acceleration (top) and AHEAD at $8\times$ acceleration (bottom).

For each dataset, the first row shows the reference image and reconstructions; the second row shows the sampling pattern in the first column and reconstruction error maps in the remaining columns; the third row shows reference and predicted segmentation overlays. Annotated scores report the dataset-level results. PSNR is reported in dB, while SSIM, Dice, and mIoU are reported as percentages. PC, FC, TC, and MEN denote patellar cartilage, femoral cartilage, tibial cartilage, and meniscus. GM, WM, and CSF denote gray matter, white matter, and cerebrospinal fluid.