



Semi-supervision for clinical contrast-weighted image synthesis from magnetic resonance fingerprinting

Mahmut Yurt¹ · Cagan Alkan¹ · Xiaozhi Cao² · Congyu Liao³ · Zihan Zhou² · Tolga Cukur⁴ · Ali Syed² · John Pauly¹ · Shreyas Vasanawala² · Kawin Setsompop^{1,2}

Received: 23 December 2025 / Revised: 11 February 2026 / Accepted: 22 February 2026
© The Author(s) 2026

Abstract

Purpose: To facilitate ease of data compilation across diverse populations for training models to synthesize clinical contrast-weighted images from magnetic resonance fingerprinting.

Methods: We leverage a semi-supervised training framework using highly accelerated acquisitions of the target contrasts used as ground truths. We utilize complementary randomized data sampling masks across training subjects and contrasts for homogeneous learning in k-space, together with multi-task learning.

Results: Our experiments indicate that the proposed method achieves high-quality synthesis with networks trained on retrospectively and prospectively undersampled data of the contrast-weighted images, enabling undersampling up to 12–16×.

Conclusions: The proposed method enables semi-supervised learning for synthesis from MRF with an end-to-end, ultra-fast training data acquisition protocol that is easier to obtain across a large population in clinical settings.

Keywords Magnetic resonance fingerprinting · Synthesis · Clinical contrasts · Semi-supervision · MRI

Introduction

Magnetic resonance fingerprinting (MRF) is a powerful quantitative imaging modality that enables simultaneous estimation of multiple tissue parameters through a single, time-efficient acquisition [1–3]. MRF leverages a tailored pulse sequence that deliberately alters scan parameters over time [1, 4]. As a result, voxels with distinct tissue characteristics produce unique signal evolutions, akin to fingerprints, when exposed to the MRF sequence [1]. The acquired fingerprints are matched against a precalculated dictionary of simulated signal evolutions, utilizing maximum correlation matching or principal component analysis (PCA) techniques [5–7]. The matching process allows for the quantification of tissue parameters, such as T_1 , T_2 , M_0 , and B_0 , resulting in

a comprehensive set of quantitative maps for the underlying anatomy [3, 8].

Recent advances in MRF driven by sequence design innovations [9, 10] and image reconstruction techniques [5, 11–13] have enabled high-resolution, whole-brain imaging protocols to be completed within just a few minutes of data acquisition. This represents a major improvement in both efficiency and quality of the imaging, offering faster and more comprehensive quantitative characterization capabilities [14–17]. However, clinical practice still predominantly relies on traditional MRI contrasts, e.g., T_1 -weighted, T_2 -weighted, and fluid attenuated inversion recovery (FLAIR) images, because of their longstanding familiarity and widespread use among radiologists [18]. Unlike the MRF acquisition which can be acquired across the whole brain at high isotropic resolution in a few minutes [11], the clinical contrast-weighted exams are prolonged, often taking up to 20–40 min [19], depending on the number of clinical contrasts involved. To facilitate the adoption of fast MRF protocols in routine diagnostics, there is a significant opportunity to translate MRF data into familiar contrast-weighted images via synthesis (input: MRF, output: contrast-weighted images), making the transition smoother for clinical use [19–22].

✉ Mahmut Yurt
myurt@stanford.edu

¹ Electrical Engineering, Stanford University, Stanford, CA, USA

² Radiology, Stanford University, Stanford, CA, USA

³ Radiology, University of California, San Francisco, San Francisco, CA, USA

⁴ Electrical Engineering, Bilkent University, Ankara, Turkey

A natural approach to synthesizing contrast-weighted images from MRF acquisitions is to utilize MR physics principles based on Bloch equations [23] or extended phase graphs (EPG) [24]. This physics-based synthesis takes the parametric information from the quantitative maps as input and simulates the desired, target contrasts [25, 26]. However, this synthesis is typically hindered by incomplete modeling of the acquisition process, shown by previous studies [25, 27–29]. In particular, addressing partial volume effects and accurately modeling multi-tissue compartments present important challenges [30, 31]. Furthermore, precisely estimating the impact of magnetization transfer that significantly affects the acquired contrast remains challenging, especially for FLAIR images [32, 33]. Consequently, the synthesized images may suffer from undesired noise artifacts, false contrasts and flow effects as well as inaccurate tissue edges [18, 33–35].

To overcome the limitations of physics synthesis, several studies have explored data-driven models utilizing deep learning techniques for synthetic MRF [18, 19, 36, 37]. For instance, Virtue et al. [36] introduced a deep artifact noise model, PixelNet, that voxel-wise maps noisy and aliased raw spatiotemporal acquisitions onto contrast-weighted images. Building on this, Wang et al. [18] proposed a multi-branch conditional generative adversarial network (GAN) to simultaneously synthesize multiple target contrasts again from raw spatiotemporal MRF acquisitions. Nykanen et al. [37] further developed a deep synthesis model by proposing a multiple U-Net based architecture to synthesize clinical contrasts from MRF parametric maps, taking advantage of dimensionality-reduced input images for deep learning. Note that the dictionary fitting process in MRF parametric map estimation will result in quantitative maps that do not contain the full information of the acquired data, as they do not capture unmodeled information, including partial volume and multi-compartment effects [33, 38]. Therefore, Schauman et al. [19] developed a deep GAN model that inputs the low-rank coefficient images from the subspace reconstruction of the MRF acquisitions, which provide a high level of dimensionality reduction while retaining most of the time-series signal evolution information contained in the raw data [33]. While all these deep learning models consistently demonstrate superior performance over physics-based approaches, they heavily depend on full supervision from high-quality contrast-weighted images as ground truth labels during training. Consequently, they require a training dataset of paired, high-resolution MRF and clinical contrast data, which should be obtained from a large and diverse clinical population with varying pathologies. However, the clinical contrasts can take a long time to acquire, given that the acceleration ratio is a limiting factor in parallel imaging, compressed sensing, or machine learning recovery of these contrasts in clinical settings [19]. Moreover, conventional clinical scans are

typically conducted at anisotropic resolutions, further complicating the compilation of high-quality, spatially paired data for training.

In recent years, semi-supervised learning (SSL) has gained significant traction in machine learning [39–42], particularly in scenarios where obtaining comprehensive, high-quality ground truth labels is costly or impractical [43–45]. The SSL framework involves training deep models with datasets that have incomplete or missing ground truth labels. By reducing the reliance on extensive labeled datasets, SSL not only captures the underlying data distributions more effectively but also results in more robust and generalizable models. This approach has been successfully applied across various domains, including natural language processing [46–49], computer vision [50–53], and medical image analysis [54–56] as well as MRI [43, 45, 57], where the scarcity of labeled data poses a significant challenge in model training. In the context of MRF to contrast-weighted image synthesis, semi-supervision can refer to a training protocol with missing k-space samples in the target ground truths of clinical contrasts. Therefore, the SSL framework can leverage highly undersampled contrast-weighted images instead of fully sampled ones as ground truth labels, thereby reducing the required scan time for training dataset compilation. Consequently, this approach should enable us to overcome the limitations of previous fully supervised methods.

Here, we propose a semi-supervised synthesis model, named ssMRF, for contrast-weighted image synthesis from MRF.¹ To reduce training data requirements, the proposed ssMRF model enables a training framework with highly accelerated acquisitions (acceleration factor, R , up to $16x$) of the clinical contrasts. Contrary to previous fully supervised models that define the loss function on the entire k-space of the target contrasts, ssMRF leverages a semi-supervised loss defined only on the acquired k-space samples of the accelerated acquisitions. The semi-supervised loss function is performed via a physics module that mimics the coil projection and undersampling of the accelerated ground truths. For homogenous learning across the k-space, ssMRF further utilizes complementary, randomized sampling masks across training subjects and contrasts, via multi-task learning to synergistically synthesize multiple contrasts. Experiments were performed on retrospectively and prospectively undersampled data. For prospective experiments, we developed undersampled acquisitions based on GE's fast turbo spin echo sequences, called CUBE, at $R = 16x$ for three clinical contrasts (T1-CUBE, T2-CUBE, FLAIR-CUBE), all at 1 mm isotropic resolution, with a total of 3 min and 42 s scan duration. In these sequences, the k-space sampling schemes were

¹ A preliminary version of the work as a 1-page abstract has been presented in the ISMRM 2023 and 2024 conferences. Abstract IDs: 0423 [58] and 3562 [59], respectively.

modified to provide uniform random sampling masks that are complementarily generated for each contrast on the fly during data acquisition, while the k-space acquisition orderings are designed to ensure that correct clinical contrasts are obtained. Our demonstrations indicate that the proposed ssMRF method can effectively synthesize target contrasts while maintaining a similar synthesis quality with the fully supervised approach.

Contributions

- To best of our knowledge, this is the first semi-supervised model for synthesizing contrast-weighted images from MRF without needing fully sampled ground truths.
- We introduce a physics module to define the semi-supervised loss on acquired k-space samples of the undersampled contrast data.
- We implement contrast scans with complementary random undersampling masks generated on the fly.
- We perform experiments on retrospectively and prospectively undersampled data up to 16x undersampling, and demonstrate near performance with the fully supervised model.

Theory

This section provides the details of the proposed framework, organized as follows: (1) subspace reconstruction for deriving coefficient images from MRF acquisitions, (2) the semi-supervised learning approach central to ssMRF, (3) multi-task learning for the simultaneous synthesis of multiple clinical contrasts, and (4) the complementary, randomized sampling masks for data collection.

Subspace reconstruction

Time-series acquisitions in MR fingerprinting track a low-rank signal evolution over time that demonstrates a strong temporal correlation [5]. In order to capture this correlation and mitigate redundancy, MRF acquisition can be represented with a set of low-rank subspace coefficient images [11], denoted as C , which is calculated through subspace reconstruction (see Fig. 1a for details). Unlike quantitative parameter maps (e.g., T_1/T_2) that compress the acquisition into a few scalars, the subspace coefficient images provide a compact low-dimensional representation of the full temporal signal evolution, which can provide additional information for downstream synthesis.

In subspace reconstruction, the MRF dictionary is first calculated using the extended phase graph technique based on the utilized flip angle (FA) and repetition time (TR) values [24]. The resulting dictionary is compressed via singular

value decomposition (SVD), and the first n temporal principal components that contain most of the information are extracted to form subspace bases Φ_{1-n} . The subspace bases are then used to recover the corresponding first n coefficient images C_{1-n} . This spatiotemporal subspace reconstruction is formulated as an optimization problem, and leveraged with a regularization term:

$$\min_{C_{1-n}} \|\mathcal{N}S\Phi C_{1-n} - x\|_2^2 + \lambda L(C_{1-n}) \quad (1)$$

where $C_{1-n} = \{c_1, c_2, \dots, c_n\}$ denotes the first n coefficient images, \mathcal{N} is the non-uniform Fourier transform defined on the undersampling pattern, S denotes the coil sensitivity maps, x is the acquired MRF data, L is the regularization term, and λ adjusts the weighing of the regularization. While the subspace reconstructed coefficient images C_{1-n} are dimensionality-reduced representations, MRF time-resolved images can still be recovered from these coefficient images for further analysis [11].

Crucially, the subspace representation helps mitigate partial volume effects. Because the projection onto the subspace basis Φ is a linear operation, the coefficient vector C for a voxel containing mixed tissues effectively represents a weighted linear combination of the individual tissue coefficients [60]. As a result, the representation can retain multi-compartment signal contributions that are often obscured when enforcing a single-compartment quantitative fit (e.g., a single scalar per voxel), providing a richer and more robust feature set as an input to the synthesis network.

Semi-supervised learning

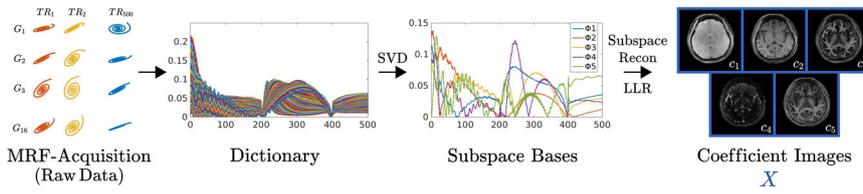
The deep learning-based synthesis of targeted clinical contrast-weighted images begins by using the MRF coefficient images as input. The synthesis problem is formulated as:

$$G(C_{1-n}) = \hat{Y} = \{\hat{y}_1, \hat{y}_2, \dots, \hat{y}_m\} \quad (2)$$

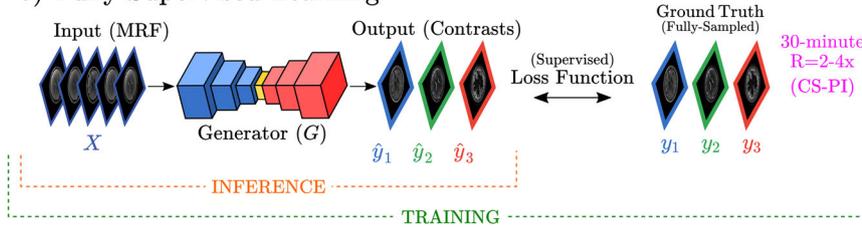
where G is the synthesizer network, C_{1-n} represents the first n MRF coefficient images used as input, and $\hat{Y} = \{\hat{y}_1, \hat{y}_2, \dots, \hat{y}_m\}$ denotes the output images as predictions of the synthesizer network for the target clinical contrasts. Here, each \hat{y}_j with $j \in \{1, 2, \dots, m\}$ corresponds to the image of a different clinical contrast.

In traditional fully supervised frameworks (see Fig. 1b), training the synthesizer network G requires high-quality, high-resolution target contrast-weighted images to be used as ground truth labels. These high-quality ground truths are typically acquired through prolonged procedures in clinical settings (e.g., 20 to 40 min acquisition times, varying based on the number of contrasts involved), and are recovered using parallel imaging, compressed sensing, or machine

a) Spatiotemporal Subspace-Reconstruction



b) Fully-Supervised Learning



c) Semi-Supervised Learning

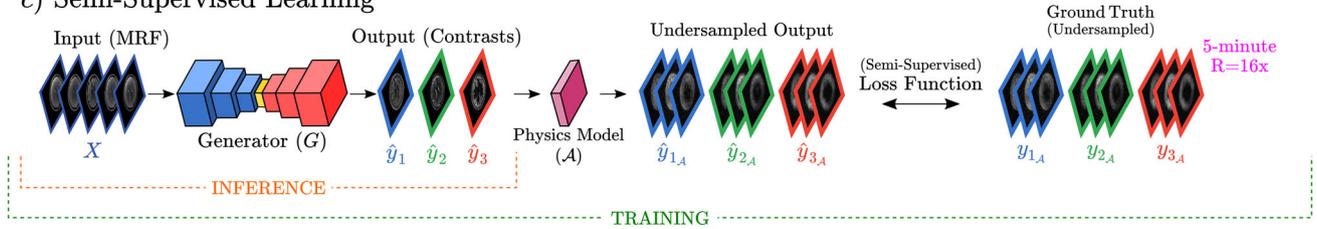


Fig. 1 a) MRF coefficient images are recovered from the raw acquisitions via subspace reconstruction. b) The prior fully supervised models rely on high-quality contrast data (e.g., 30 min, $R = 2-4x$, parallel imaging reconstructed) for direct loss computation. c) The proposed semi-supervised learning framework allows model training with highly

accelerated acquisitions of the contrast-weighted images as ground truth labels (e.g., 5 min, $R = 16x$). It incorporates a physics module to generate undersampled multi-coil output to define the semi-supervised loss function only on acquired k-space points

learning techniques with moderately low acceleration factors R around $2-4x$ [19]. However, the extensive data collection times necessary for clinical contrasts make it impractical to compile a comprehensive dataset of paired MRF and contrast-weighted images, especially across diverse clinical populations needed for robust training.

To reduce data requirements, the proposed semi-supervised approach, named ssMRF (see Fig. 1c), enables ground truth contrast-weighted labels to be highly accelerated acquisitions (e.g., 5-min acquisition total across 4 target contrasts at high isotropic resolution whole-brain, $R \simeq 16x$), denoted as Y_{R_x} . Since the accelerated acquisitions Y_{R_x} suffer from severe aliasing artifacts, it is infeasible to apply fully supervised loss functions between the synthesized images and accelerated acquisitions of the target contrasts. Therefore, ssMRF introduce a physics-based module \mathcal{A} that generates accelerated multi-coil counterparts of the synthesized images:

$$\hat{Y}_{R_x} = \mathcal{A}(G(C_{1-n})) = \mathcal{A}(\hat{Y}) = \mathcal{F}^{-1} P_Y \mathcal{F} S_Y(\hat{Y}) \quad (3)$$

where \hat{Y}_{R_x} denotes the multi-coil, accelerated counterparts of the synthesized images, \mathcal{A} is the physics guidance module, \mathcal{F}

is the forward Fourier transform, \mathcal{F}^{-1} is the inverse Fourier transform, P_Y denotes undersampling masks, S_Y denotes coil sensitivity maps, G is the synthesizer network, and C_{1-n} denotes the input MRF coefficients. Having the synthesized undersampled counterparts \hat{Y}_{R_x} , the semi-supervised loss function is defined selectively on acquired k-space points of the ground truth contrasts (i.e., on the k-space coefficients within sampling masks of the contrasts). The overall loss function contains three distinct components: k-space, image, and adversarial losses.

k-space Loss: A Fourier domain loss function between the synthesized and ground truth target contrast-weighted images is defined on acquired k-space data:

$$L_{\text{ksp}} = \mathbb{E}_{C_{1-n}, Y_{R_x}} [||\mathcal{F}\mathcal{A}(G(C_{1-n})) - \mathcal{F}Y_{R_x}||_1] \quad (4)$$

Image Loss: The robustness of the synthesizer is further enhanced with a loss function between image-domain data of the synthesized and ground truth acquisitions:

$$L_{\text{img}} = \mathbb{E}_{C_{1-n}, Y_{R_x}} [||\mathcal{A}(G(C_{1-n})) - Y_{R_x}||_1] \quad (5)$$

Adversarial Loss: The level of realism in the synthesized images is enhanced with an adversarial loss function through a discriminator module D based on synthesized and ground truth acquisitions:

$$L_{\text{adv}} = -\mathbb{E}_{Y_{R_x}} [(D(Y_{R_x}) - 1)^2] - \mathbb{E}_{C_{1-n}} [D(\mathcal{A}(G(C_{1-n})))] \quad (6)$$

These sub loss functions are combined to form the final objective in a min–max game between the synthesizer and discriminator:

$$\min_G \max_D \{\lambda_{\text{ksp}} L_{\text{ksp}} + \lambda_{\text{img}} L_{\text{img}} + \lambda_{\text{adv}} L_{\text{adv}}\} \quad (7)$$

where λ_{ksp} , λ_{img} , and λ_{adv} are the weighting of the k-space, image, and adversarial losses, respectively.

Multi-task learning for multiple target contrasts

Multi-task learning (MTL) is an important approach in deep learning that trains multiple related tasks simultaneously, leveraging shared representations to improve the performance of each individual task [61, 62]. By sharing information across tasks, MTL can produce more robust and generalizable models, while also reducing the need for extensive training datasets for each specific task. In the context of MRF contrast synthesis, MTL can be employed to generate multiple target contrast-weighted images, exploiting the inherent commonalities between the contrasts. The multi-task synthesizer and discriminator networks can be summarized as follows:

Multi-task synthesizer: The synthesizer network G composes both shared G_S and task-specific G_T layers. The shared layers G_S are designed to capture the shared features from the input MRF coefficient images, resulting in unified representation that can be utilized for generating different contrasts. Mathematically, this can be expressed as:

$$Z = G_S(C_{1-n}) \quad (8)$$

where Z is the shared feature representation.

Continuing from the shared feature representation, the task-specific layers G_T take the shared features Z and further process them to generate the target contrast-weighted images. These layers are tailored to capture the unique characteristics required for each specific contrast. The output of the multi-task synthesizer can be expressed as:

$$y_i = G_{T_i}(Z) \quad (9)$$

where \hat{y}_i is the generated image and G_{T_i} is the task-specific layers for the i th contrast.

Multi-task discriminator: The discriminator network D is composed solely of task-specific layers, resulting in individual discriminators D_i for the i th contrast. Each D_i is designed to distinguish between real and generated images for the corresponding contrast.

Complementary sampling masks

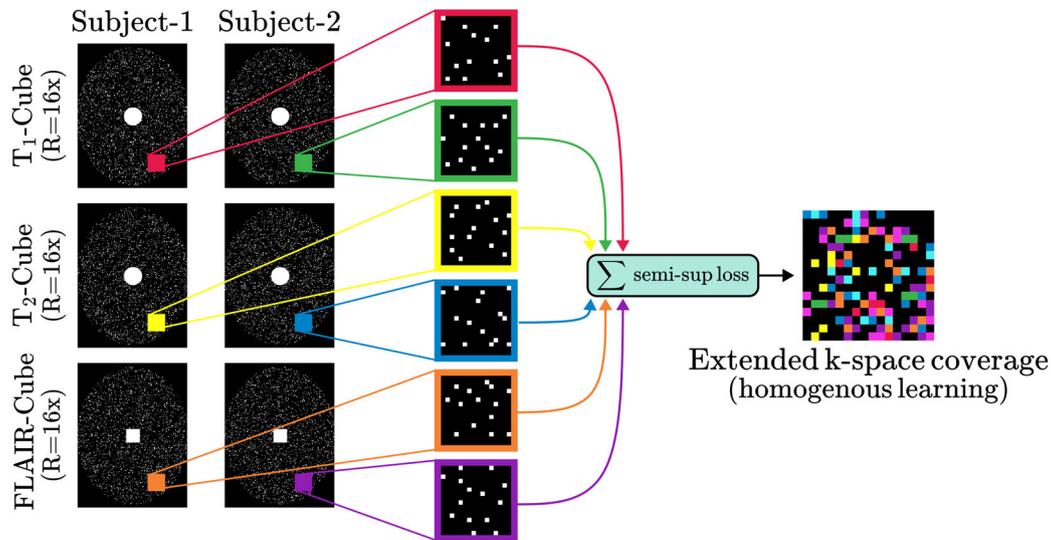
Note that the semi-supervised loss function is defined only on the acquired k-space data of the contrast acquisitions, ensuring that backpropagation in deep learning only occurs through these points. Consequently, using the same sampling mask for all training subjects and contrasts would restrict the network to learning to recover only the k-space points within that specific sampling mask. To prevent this limitation, complementary randomized sampling masks are utilized across the subjects and contrasts. These masks allow the semi-supervised loss function to span all possible k-space locations when summed across subjects and contrasts. Here we demonstrate two possible mask types that can be used for complementary sampling (please see Fig. 2). The first type is the complementary uniform random (CUR) masks (Fig. 2a), where the masks are generated using a uniform random distribution in accordance with the acceleration ratio R . The second type is the complementary Poisson disk (CPD) masks (Fig. 2b), where the masks are generated using a variable-density Poisson disk algorithm that ensures a minimum distance between sampled points, thereby reducing the occurrence of large gaps in k-space coverage. The CPD masks are generated by iteratively selecting random points in k-space, ensuring that each selected point maintains a minimum distance from previously selected points across subjects and contrasts.

Methods

Dataset

Retrospectively undersampled data: We collected an in-house dataset of MRF and four clinical contrasts (T₁-MPRAGE, T₁-CUBE, T₂-CUBE, and FLAIR-CUBE) using a 3T GE Premier scanner from 32 subjects at the Lucas Center at Stanford University. All data were collected with institutional review board (IRB) approval, and written informed consent was obtained from all subjects. For all sequences, data collection was performed at 1 mm isotropic resolution across the whole brain. The MRF data were acquired using the 3D tiny-golden-angle-shuffling multi-axis spiral-projection-imaging (3D-TGAS-SPI) protocol [11], consisting of 48 groups, resulting in a total acquisition time of 5:57 min per subject. The contrast-weighted images were acquired using clinically employed MPRAGE and CUBE

a) Complementary Uniform Random (CUR) Masks



b) Complementary Poisson Disc (CPD) Masks

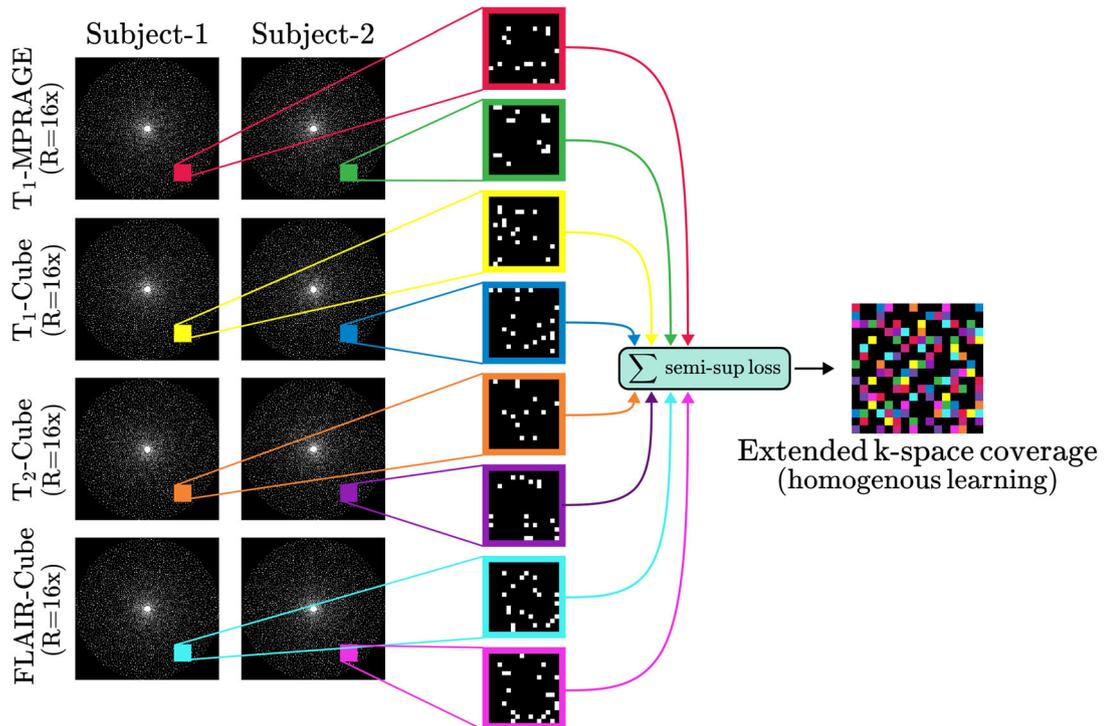


Fig. 2 **a)** Complementary uniform random (CUR) masks are displayed for three clinical contrasts and two subjects. **b)** Complementary Poisson disc (CPD) masks are displayed for four clinical contrasts and two subjects. Both of the complementary mask types involve a random-

ization in the sampling masks across the contrasts and subjects that helps extend k-space coverage for the semi-supervised loss function. Zoomed-in regions from the masks are displayed together with their combination yielding the extended k-space coverage

sequences with acquisition times of approximately 4:57 min for T_1 -MPRAGE at $R = 2x$, 5:05 min for T_1 -CUBE at $R = 3x$, 3:22 min for T_2 -CUBE at $R = 4x$, and 6:21 min for FLAIR-

CUBE at $R = 4x$. These lightly accelerated sequences resulted in a total acquisition time of ~ 20 min, which was deemed reasonable to perform across our test subject group, while a

nonaccelerated acquisition protocol would have taken around one and a half hours to perform, which would have been very lengthy and likely prone to introduce motion artifacts. The raw k-space data were recovered from the lightly accelerated protocol with parallel imaging and compressed sensing reconstruction using the GE Orchestra Software Development Kit and Berkeley Advanced Reconstruction Toolbox (BART) [63]. For experiments, the reconstructed k-space data were treated as fully sampled acquisitions and then they were retrospectively undersampled to generate training data for ssMRF as a proof-of-concept demonstration. Among the acquired 32 subjects, we reserved 24 subjects for training, 3 for validation and 5 for testing.

Prospectively undersampled data: We then collected a dataset with prospective undersampling across 20 additional volunteers, where MRF and three clinical contrasts (T_1 -CUBE, T_2 -CUBE, and FLAIR-CUBE) were acquired using the same 3T GE Premier scanner. The MRF acquisition remains the same as before while the contrast-weighted acquisitions were directly obtained from highly accelerated sequences all at 16x acceleration that we implemented on the scanner with complementary uniform random (CUR) sampling based on GE's fast spin echo CUBE sequences. These sequences corresponded to a scan duration of 1:23 min for T_1 -CUBE, 0:55 min for T_2 -CUBE, and 1:24 min for FLAIR-CUBE. To test the performance of the model trained with the prospectively undersampled data, we additionally collected 2 more subjects which we used only for testing. For these testing subjects, we collected 5:57 MRF and three clinical contrasts T_1 -CUBE 5:56min, T_2 -CUBE 3:58min, FLAIR-CUBE 5:06min; each at $R = 3x$). The contrast-weighted acquisitions of the test subjects were reconstructed with L1-wavelet and served as reference in qualitative and quantitative evaluations during testing.

Implementation details

Subspace reconstruction: The number of subspace coefficient images, denoted as n in Sect. 2.1, was set to 5 based on both quantitative and visual assessments. Quantitatively, analysis on the singular value spectrum of the dictionary reported in Supp. Table S1 reveals that the first five principal components capture 99.92% of the signal energy, with additional components yielding negligible gain. Visually, inspection of the subspace coefficients shown in Supp. Fig. S1 confirmed that distinct structural information is concentrated within the first five components, while higher-order coefficients are dominated by noise. Meanwhile, locally low-rank regularization was set as the regularization term L , with the weighting parameter λ optimized to be 5×10^{-3} .

Semi-supervised synthesis and multi-task learning: The semi-supervised synthesis network is designed to gener-

ate target contrast-weighted images, \hat{Y} . Its input comprises coefficient images C_{1-5} obtained from subspace reconstruction, where the magnitude component was provided to the network. Here, we employed a dual-generator framework for training: one generator predicts the magnitude (absolute value) images of the target contrasts, and the other predicts the corresponding phase maps of the target contrasts. Although only the magnitude images are required as the final output, accurate phase estimation is essential during training to define the semi-supervised loss. Specifically, the predicted magnitude and phase maps are combined to form complex-valued images of the target contrasts, which are subsequently passed through the forward physics model (accounting for coil sensitivities and undersampling) and compared against the ground-truth multi-coil, undersampled k-space data. To facilitate this estimation, the phase generator is additionally conditioned on phase maps obtained from zero-filled reconstructions of the undersampled ground-truth data of the target contrasts. These phase maps are used directly as input channels, without any additional processing or smoothing, and are available only during training. In contrast, the magnitude generator is supplied only with the coefficient images. Consequently, at inference time, the phase generator is discarded, and only the magnitude generator is utilized to synthesize the final output images. For both phase and magnitude generators, the multi-task synthesis network G consisted of two components: a shared module and contrast-specific modules. The shared module, used across all target contrasts, consisted of an encoder with three convolutional layers followed by eight residual (ResNet) blocks. Meanwhile, each output contrast had its own dedicated module comprising a single ResNet block and a decoder with three convolutional layers. A contrast-specific discriminator network D of five convolutional layers was also implemented for each output contrast. The training was conducted over a total of 100 epochs. The learning rate was initialized at 0.0002 for the first 50 epochs and then linearly decayed to 0 over the remaining 50 epochs. The relative weights for the k-space, image, and adversarial domain loss functions were selected based on validation performance.

Complementary random sampling masks: For retrospectively undersampled experiments, we designed complementary Poisson disk (CPD) masks with acceleration ratios of $R = 4x$, $R = 8x$, and $R = 12x$. These masks utilize a variable-density algorithm to maintain minimum spacing between samples. For prospectively undersampled data, complementary uniform random (CUR) masks were utilized with an acceleration ratio of $R = 16x$. In both cases, the masks are complementary in that they are randomized distinctively across subjects and contrasts to maximize k-space coverage for the semi-supervised loss function. All masks included a fully sampled

calibration region of size 12×12 to condition the learning problem with low-frequency information.

View ordering and contrast preservation: To accommodate the signal evolution associated with long echo trains, we employed a standard flexible view-ordering algorithm routinely used in clinical practice (please see Fig. 3 for visual demonstrations). This algorithm sorts the acquisition order based on k-space coordinates such that successive samples are acquired from proximate k-space locations, yielding a smooth modulation transfer function and reducing image artifacts.

For T₁-CUBE, we utilized a radial modulation view-ordering approach. In this approach, phase-encode views (k_y, k_z) are sorted by their radial distance (d_r) from the k-space origin. Consequently, the earliest echoes in the train, which exhibit the highest signal amplitude, are mapped directly to the center of the k-space ($d_r \simeq 0$). This mapping minimizes the effective echo time preserving the intended T₁-weighted contrast generated at the onset of the echo train.

For T₂-CUBE and FLAIR-CUBE, we employed a circular modulation strategy. In this approach, views are sorted along the principal phase-encode axis to explicitly control the effective echo time by assigning the center of the k-space to the central echo index. By aligning the central k-space acquisition with echoes occurring midway through the train, the sequence captures the steady-state magnetization reflecting the desired T₂ preparation or inversion recovery contrast.

Evaluation metrics

In our assessments of synthesis quality, we used two well-established MRI metrics: peak signal-to-noise ratio (PSNR) and structural similarity index measure (SSIM). More specifically, PSNR provides a quantitative measure of synthesis fidelity by assessing the peak signal strength to mean squared error, defined in decibel (dB) as follows:

$$\text{PSNR} = 20 \cdot \log_{10} \left(\frac{\text{MAX}_I}{\text{MSE}} \right) \quad (10)$$

where I is the reference image, MAX_I denotes the maximum pixel intensity of image I , and MSE is mean squared error which is formulated as follows:

$$\text{MSE} = \frac{1}{mn} \sum_{i=0}^{n-1} \sum_{j=0}^{m-1} [I(i, j) - \hat{I}(i, j)]^2 \quad (11)$$

where I is the reference image, \hat{I} is the synthesized image.

Furthermore, we leveraged SSIM to complement our evaluation by considering perceptual aspects and structural

similarity between reference and synthesized images:

$$\text{SSIM}(x, y) = \frac{(2\mu_x\mu_y + c_1)(2\sigma_{xy} + c_2)}{(\mu_x^2 + \mu_y^2 + c_1)(\sigma_x^2 + \sigma_y^2 + c_2)} \quad (12)$$

where x and y are two windows of common size, μ_x is the mean of x , μ_y is the mean of y , σ_x^2 is the variance of x , σ_y^2 is the variance of y , σ_{xy} is the covariance of x and y , and c_1, c_2 are two variables to stabilize the division by a weak denominator.

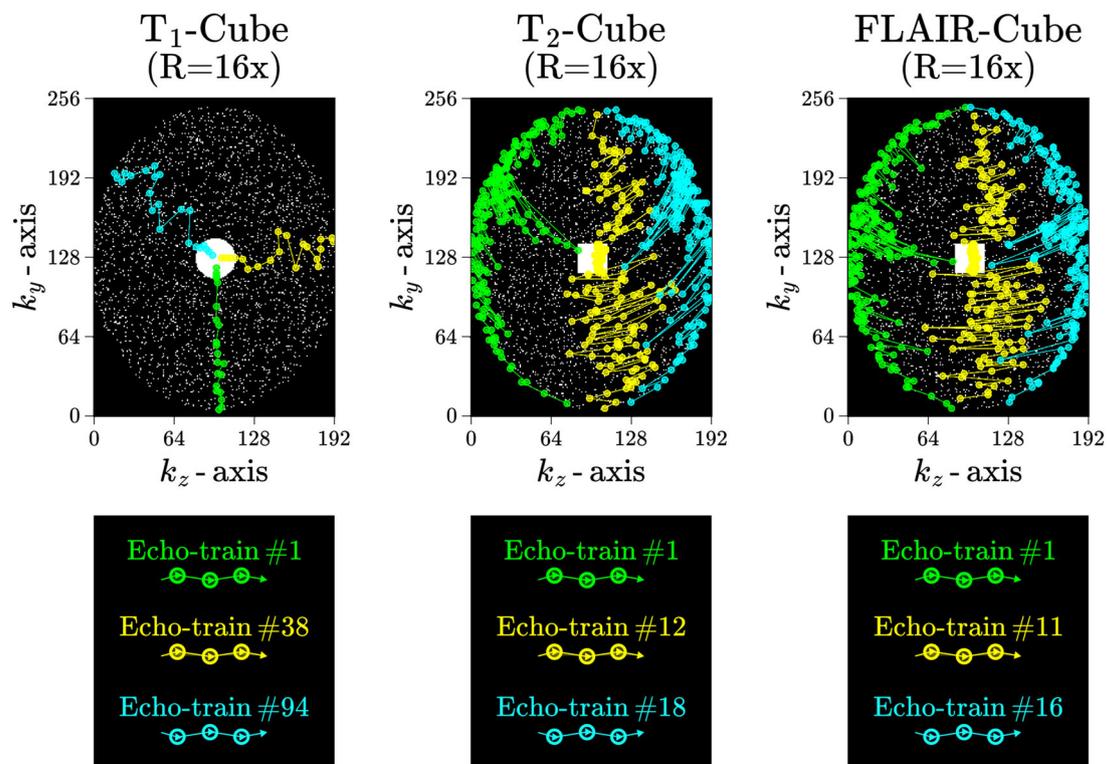
Results

Demonstrations on retrospectively undersampled data

We first demonstrated the proposed semi-supervised ssMRF method on the retrospectively undersampled data. The dataset included synthesis tasks targeting four clinical contrasts: T₁-MPRAGE, T₁-CUBE, T₂-CUBE, and FLAIR-CUBE. For demonstrations, separate ssMRF models were trained at different acceleration ratios $\{R=4x, 8x, 12x\}$, yielding the semi-supervised models ssMRF ($R=4x$), ssMRF ($R=8x$), ssMRF ($R=12x$), respectively. These models were compared against a fully supervised benchmark, fsMRF, trained on fully sampled acquisitions ($R=1x$) of the target clinical contrasts. Additionally, we included comparisons against a physics-based, non-learning synthesis approach using Bloch equation simulations.

To assess the robustness of ssMRF to acceleration ratio and training data size, we evaluated the synthesis quality across different acceleration ratios $R = \{4x, 8x, 12x\}$ and varying number of training subjects $n_T = \{8, 12, 16, 20, 24\}$. PSNR measurements between the synthesized and reference images on the test set are reported in Fig. 4 for all target contrasts. For comparison, the synthesis performance of the fully supervised fsMRF model across the number of training subjects is also reported, along with the results from Bloch equation synthesis. Reported measurements indicate that, for T₁-MPRAGE, T₁-CUBE, and T₂-CUBE, increasing the number of training subjects consistently improved performance for both ssMRF and fsMRF, while the acceleration factor has only a minimal effect when the training set size is fixed. Meanwhile, for FLAIR-Cube, increasing the number of training subjects does not yield strictly monotonic improvements, particularly at higher acceleration factors ($R=8x$ and $R=12x$). We attribute this behavior to the lower intrinsic SNR of the FLAIR acquisition, which limits the benefits of additional training data under aggressive undersampling where the supervision signal from highly accelerated data might get insufficient for ssMRF to utilize robust features. Nevertheless, for high-SNR contrasts, these results suggest

a) Ordering of Sampling Mask Points



b) Zero-Filled Images of Accelerated Contrasts

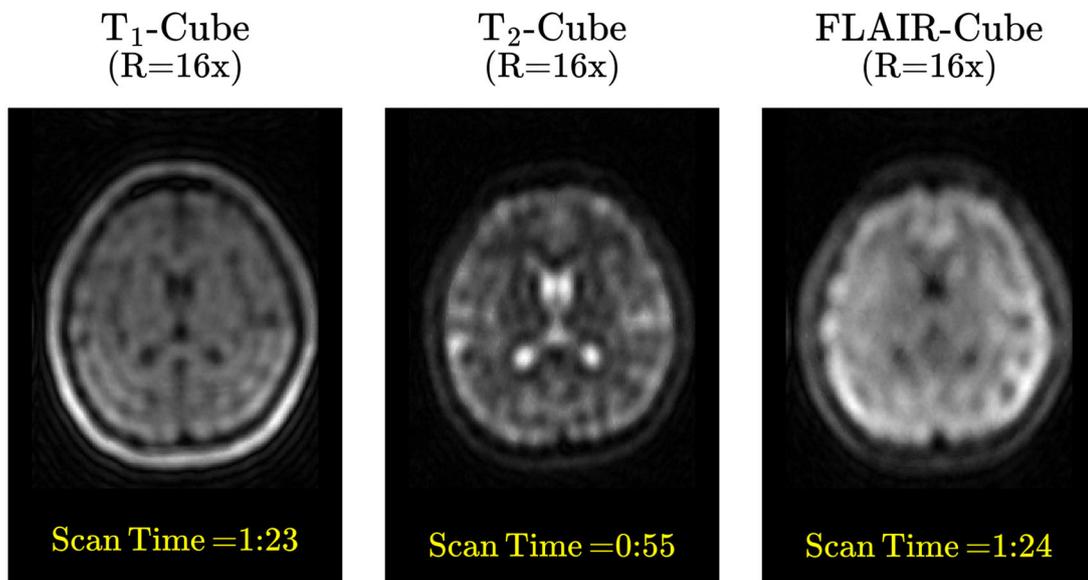


Fig. 3 a) In order to capture the correct tissue contrast obtained in standard clinical protocols, the points within the sampling masks are carefully ordered and echo-train lengths are determined. b) Zero-filled

reconstructions of the accelerated contrast-weighted acquisitions are displayed, together with corresponding scan times obtained in the prospectively undersampled sequences

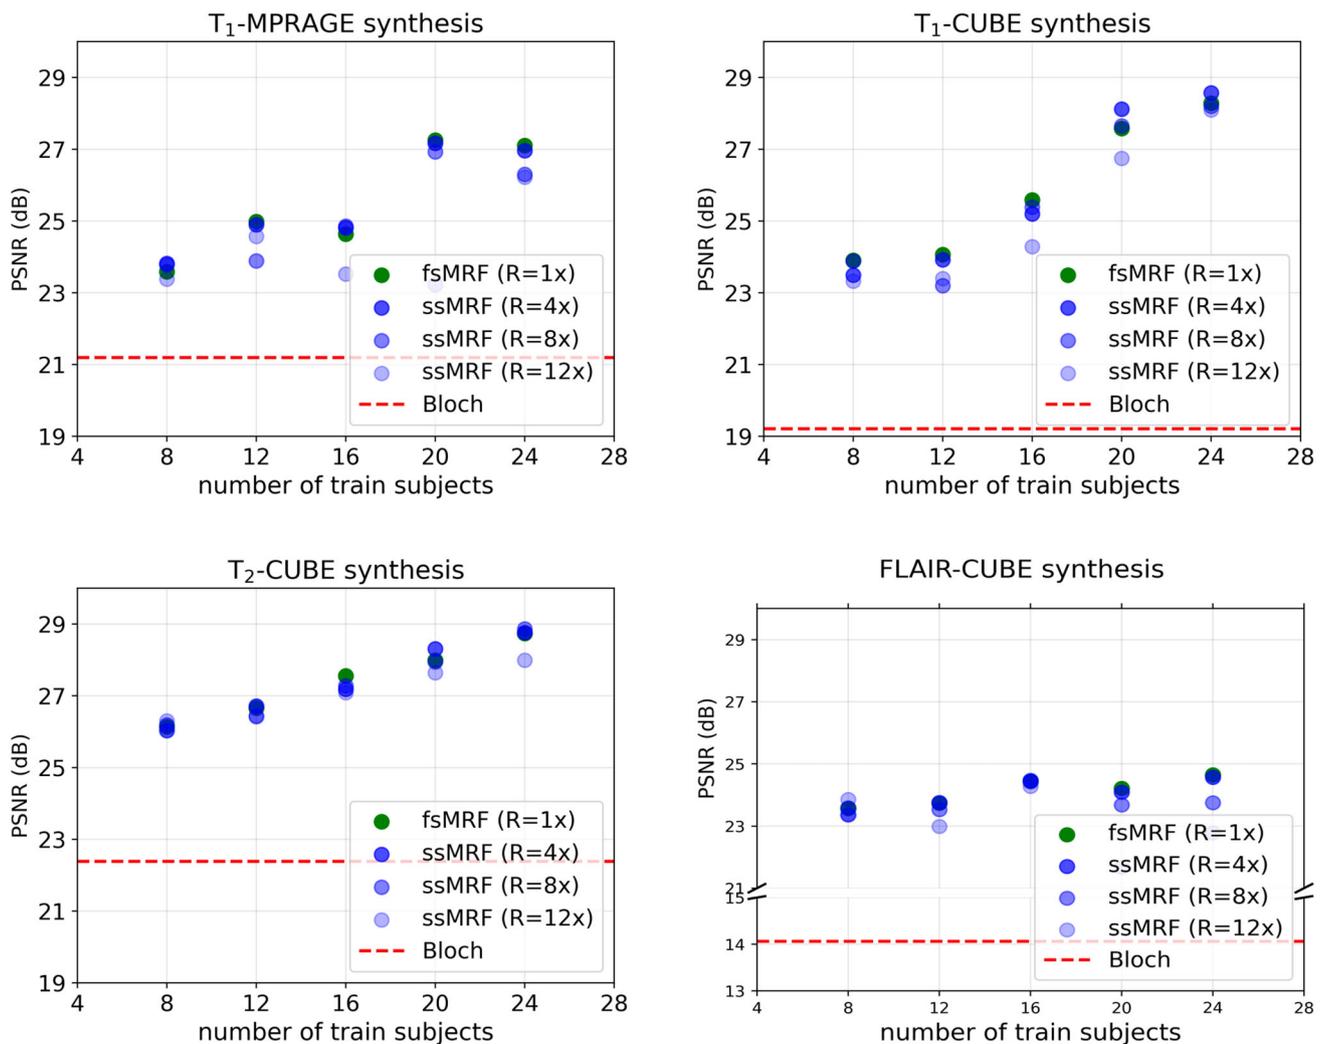


Fig. 4 PSNR measurements for the proposed ssMRF model compared to benchmarks across varying training set sizes ($n_T = 8, 12, 16, 20, 24$). Each scatter point represents the mean PSNR (dB) on the test set for four target contrasts: T₁-MPRAGE, T₁-CUBE, T₂-CUBE, and FLAIR-CUBE. Blue points indicate the ssMRF model trained at acceleration factors of $R = 4x, 8x,$ and $12x$. Green points indicate the fully supervised fsMRF model ($R = 1x$), and red dashed

lines denote the training-free Bloch equation synthesis. To facilitate consistent visual comparison, the y-axis ranges have been standardized to 19–30 dB for T₁-MPRAGE, T₁-CUBE, and T₂-CUBE. For FLAIR-CUBE, a broken y-axis (13–15 dB and 21–30 dB) is utilized to accommodate the lower reference value of the Bloch simulation while maintaining a visual scale consistent with the other contrasts

that, under a fixed total scan time, ssMRF enables the acquisition of data from a larger and more diverse subject pool by leveraging accelerated acquisitions, leading to improved synthesis quality relative to fsMRF. For example, as illustrated in Fig. 4, instead of acquiring fully sampled data for 8 subjects for fsMRF, the same scan time could accommodate data acquisition for 24 subjects using ssMRF at $R = 12x$, yielding an approximate increase of 2.5 dB in PSNR.

Next, we reported the PSNR and SSIM measurements with standard deviation intervals for fsMRF, ssMRF ($R = 4x$) and Bloch equations over the target clinical contrasts in Table 1. The reported measurements indicate that the pro-

posed ssMRF model yield similar quantitative performance with the fully supervised model despite being trained on undersampled acquisitions with missing k-space samples. Meanwhile, both ssMRF and fsMRF yield superior synthesis performance against the non-learning-based Bloch equation synthesis.

Representative synthesis results from distinct cross-break-sectional slices on the test set are displayed in Fig. 5, accompanied by corresponding error maps to show the pixel-wise synthesis deviation across methods under comparison. For a more granular assessment, zoomed-in views of different regions from different slices are provided in

Table 1 PSNR and SSIM measurements comparing ssMRF ($R = 4x$), fsMRF, and Bloch synthesis for the four clinical contrasts in the dataset

	T ₁ -MPRAGE		T ₁ -CUBE		T ₂ -CUBE		FLAIR-CUBE	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
ssMRF ($R = 4x$)	27.00	0.908	28.59	0.924	28.80	0.908	24.61	0.828
	± 1.64	± 0.018	± 1.19	± 0.011	± 1.63	± 0.021	± 2.10	± 0.051
fsMRF ($R = 1x$)	27.15	0.910	28.30	0.908	28.78	0.896	24.70	0.836
	± 1.75	± 0.019	± 0.93	± 0.008	± 1.56	± 0.021	± 1.99	± 0.062
Bloch equation	21.23	0.775	19.23	0.725	22.42	0.789	14.08	0.512
	± 0.82	± 0.032	± 1.26	± 0.036	± 1.36	± 0.032	± 0.40	± 0.041

Supp. Figs. S2 and S3 to highlight fine-grained details. The qualitative evaluation indicates that ssMRF yields visually identical performance compared to the fully supervised fsMRF. Comparisons with Bloch synthesis reveal distinct trade-offs. While Bloch synthesis appears sharper in certain high-frequency textures for T₁- and T₂-weighted contrasts, it is hindered by significant noise and artifacts, and fails to accurately reproduce the FLAIR contrast due to unmodeled magnetization transfer effects. Conversely, while the ssMRF and fsMRF outputs exhibit moderate smoothing of fine details, they achieve significantly higher accuracy and lower residual intensities, as shown in Fig. 5, and in Supp. Figs. S2 and S3. Thus, these results indicate that learning-based methods provide a more accurate approximation of the ground truth contrasts despite the reduction in texture sharpness.

Finally, we assessed the downstream performance of the proposed ssMRF model. Segmentation maps via FreeSurfer based on the synthesized MPRAGE images from the methods under comparison are shown in Fig. 6 together with the segmentation maps from the reference image. The segmentation results validate that the proposed ssMRF model yields equivalent performance compared to fsMRF in such a downstream image analysis task.

Demonstrations on prospectively undersampled data

We then extended our experiments to the prospectively undersampled data. In this dataset, we considered the synthesis task of MRF to three clinical contrast-weighted images: T₁-CUBE, T₂-CUBE, and FLAIR-CUBE. To perform the synthesis, we trained an ssMRF model with the clinical contrast acceleration ratio of $R = 16x$, yielding ssMRF ($R = 16$) model. Since this dataset contained only prospectively undersampled data, there was not any fully sampled training data. Therefore, we did not train any benchmark fsMRF model. Representative synthesis results are displayed in Fig. 7, with zoomed-in regions provided in Supp. Fig. S4 to highlight fine-grained details. The figure compares the proposed ssMRF output against two reference images: a high-quality

acquisition ($R = 3x$) and a highly accelerated acquisition ($R = 16x$). While the ssMRF model demonstrates strong contrast fidelity, we observe varying degrees of deviation in high-frequency details relative to the $R = 3x$ reference. These differences are attributable to the inherent difficulty of the training task, as the model was supervised exclusively using $R = 16x$ data. As illustrated by the $R = 16x$ reference column (shown here for a test subject to indicate the degree of undersampling), the supervision signal available during training was heavily dominated by aliasing artifacts and noise. Although this experiment employed an aggressive acceleration factor to stress-test the method, future clinical protocols could adopt slightly more conservative acceleration rates to further minimize the domain gap between synthetic and fully sampled references.

Multi-task learning

We finally performed ablation experiments to demonstrate the efficacy of the proposed multi-task learning framework. For this purpose, we compared the performance of the multi-task ssMRF model and the single-task ssMRF model. PSNR and SSIM measurements from the multi-task and single-task ssMRF models are reported in Table 2. The reported measurements indicate that the inclusion of multi-task learning consistently improve the synthesis quality, where the average increase in PSNR is 0.52 dB and the average increase in SSIM is 1.1% ($p < 0.05$, except PSNR on T₂-CUBE and FLAIR-CUBE, and SSIM on T₁-CUBE).

Discussion

This study introduces a novel semi-supervised synthesis model, ssMRF, designed to translate MRF acquisitions into clinical contrast-weighted images. Unlike previous fully-supervised models, ssMRF enables an efficient training framework using highly undersampled acquisitions of the contrast-weighted images as ground truth labels. This is achieved by leveraging a physics guidance module and a semi-supervised loss function that selectively apply the net-

Fig. 5 Representative synthesis results are presented for the proposed semi-supervised ssMRF model and the benchmark fully supervised MRF fsMRF model, alongside reference images and training-free Bloch-based synthesis. The input MRF coefficient maps are also shown. Results are provided for T₁-MPRAGE, T₁-CUBE, T₂-CUBE, and FLAIR-CUBE synthesis tasks. Overall, the proposed ssMRF model achieves performance comparable to the benchmark fsMRF model, despite being trained with 8-fold accelerated acquisitions. Moreover, both ssMRF and fsMRF models outperform the Bloch-based synthesis, particularly in FLAIR-CUBE synthesis where magnetization transfer effects significantly impact the final image contrast

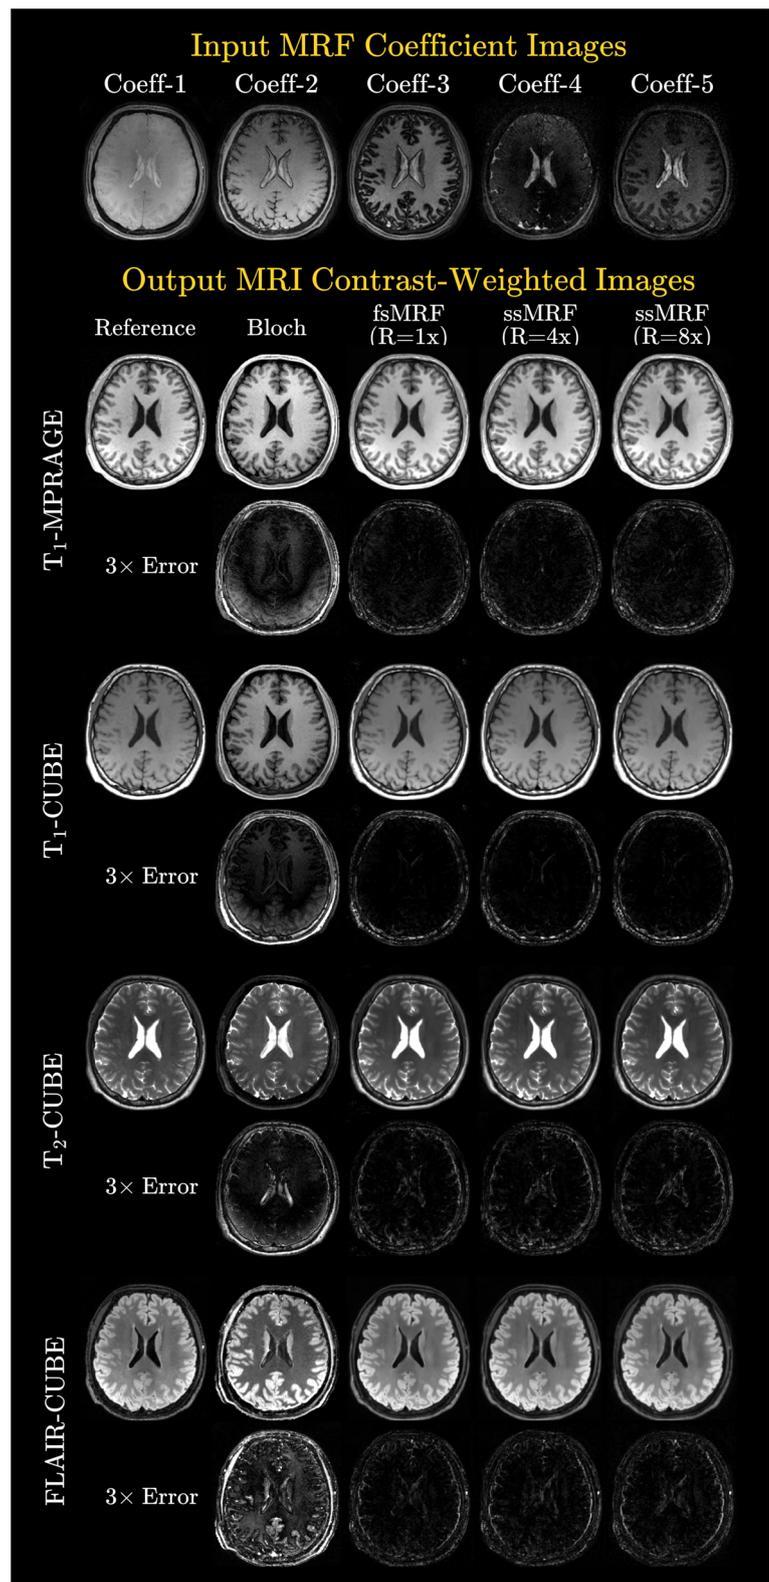
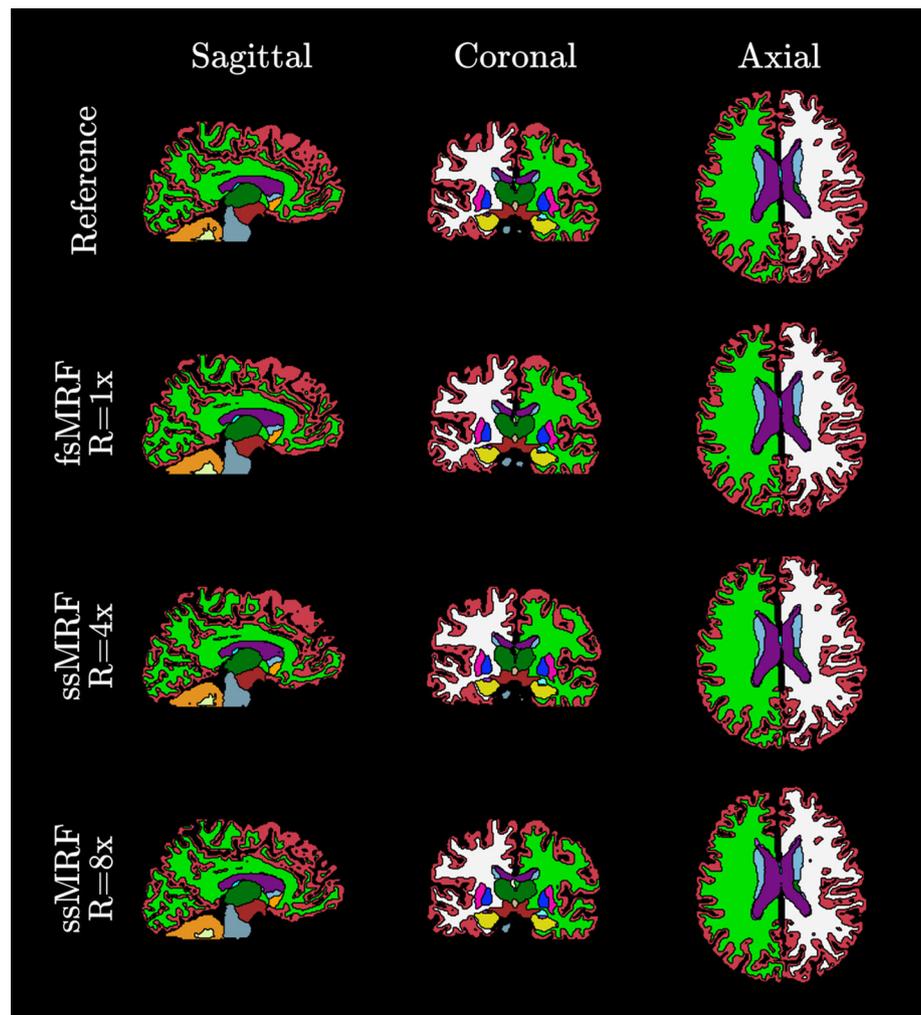


Fig. 6 Synthesized MPRAGE images from the fully-supervised and semi-supervised models were subjected to FreeSurfer segmentation together with the fully sampled reference images. Segmentation maps displayed here demonstrate that the synthesis accuracy of the semi-supervised model is sufficient enough to obtain similar segmentation maps with the fully-supervised fsMRF model as well as fully sampled references



work backpropagation only on the acquired k-space samples of the undersampled ground truths. Our experiments performed on both retrospective and prospective undersampling demonstrate that ssMRF achieves high synthesis accuracy and its performance is near with that of the fully-supervised model.

By reducing training data requirements, ssMRF facilitates the compilation of diverse datasets from larger subject cohorts, enabling the development of robust and generalizable synthesis models. In this study, we demonstrated that ssMRF models can be trained using clinical contrast ground truths accelerated up to $R = 16x$. While we achieved reasonable results utilizing only 10 subjects, increasing the training sample size significantly improves performance (see Fig. 3), allowing ssMRF to maintain parity with fully-supervised models. This trajectory suggests that with larger training sets, ssMRF could support even higher acceleration ratios. Consequently, this approach further relaxes data compilation requirements and accelerates contrast exams, enhancing the

overall efficiency and practicality of semi-supervised learning.

It is also important to consider the impact of the inherent signal-to-noise ratio (SNR) of the target clinical contrasts on semi-supervised training performance. We observed that the performance of the proposed ssMRF model improved more consistently with larger training sets and more closely matched with the fully-supervised benchmark on contrasts with higher intrinsic SNR, such as T_1 -MPRAGE, T_1 -CUBE, and T_2 -CUBE. Meanwhile, synthesizing the FLAIR contrast, which typically exhibits lower SNR, demonstrated stronger sensitivity to undersampling. Specifically, while performance degradation in our analysis was minor for other contrasts up to $R = 12x$, FLAIR exhibited a larger performance dip at this level. This suggests that performance degradation in semi-supervised learning is more pronounced in low-SNR regimes. Consequently, rather than enforcing uniform acceleration across contrasts, the undersampling budget could be redistributed based on sensitivity. High-SNR contrasts demonstrated resilience to aggressive undersam-

Fig. 7 Synthesis quality for the semi-supervised model learned on prospective undersampling is depicted. Synthesized images from the semi-supervised model are displayed together with two reference images for each contrast. $R = 3x$ reference images are reconstructed with L1-wavelet and can be considered as perfect images to evaluate synthesis quality. $R = 16x$ reference images are provided here to show the level of acceleration used to train the semi-supervised model (this is a subject from the test set, and here we did not perform any additional fine-tuning on $R = 16x$ data)

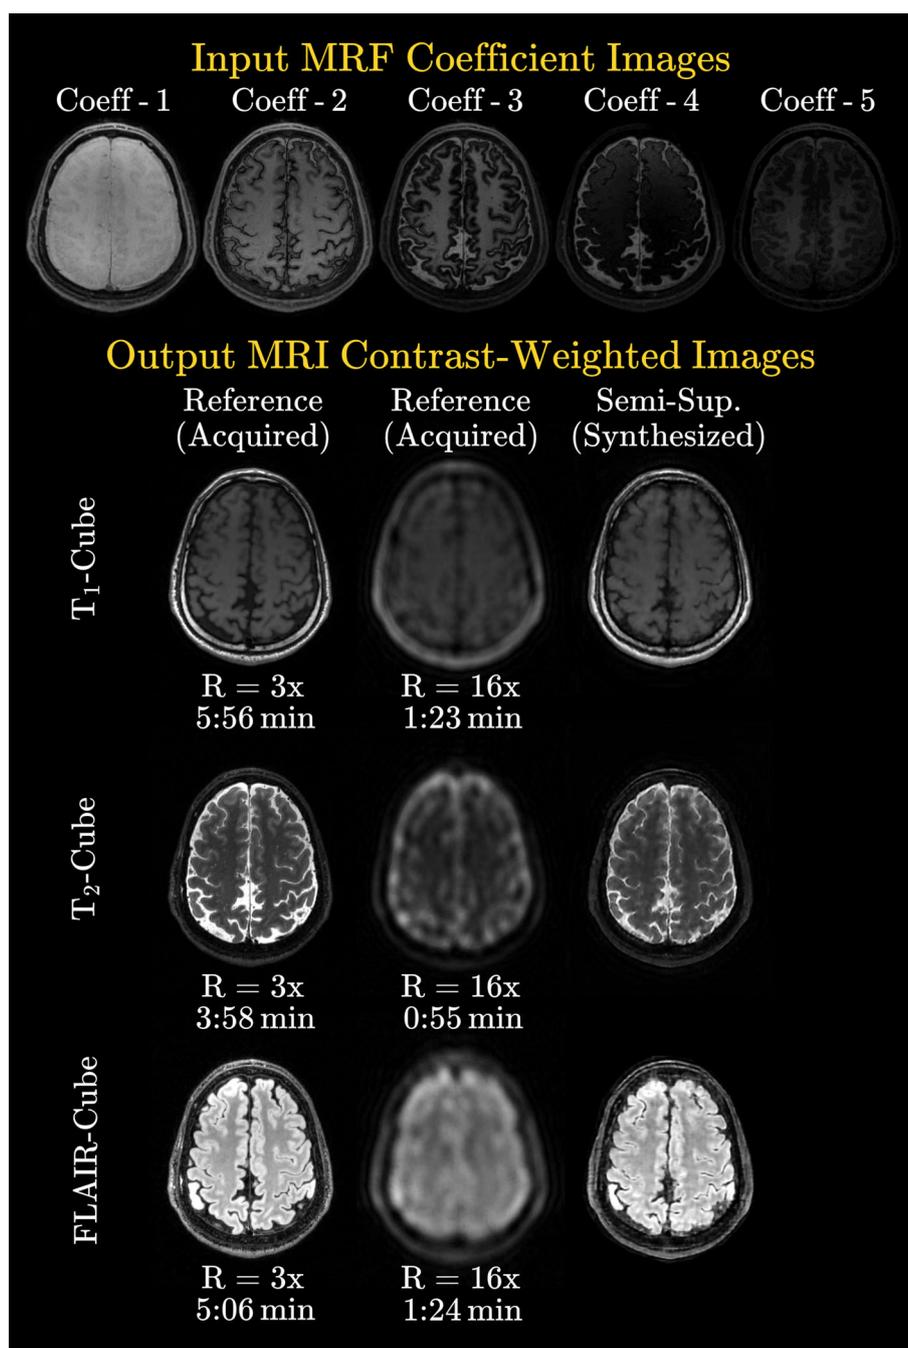


Table 2 PSNR and SSIM measurements comparing the multi-task and single-task learning frameworks that ssMRF involved

	T_1 -MPRAGE		T_1 -CUBE		T_2 -CUBE		FLAIR-CUBE	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
Multi-task	26.39	91.15	23.62	78.48	24.85	86.99	26.40	89.70
	± 1.27	± 3.65	± 1.43	± 4.56	± 1.68	± 3.44	± 1.56	± 3.56
Single-task	25.27	89.78	23.16	78.00	24.89	85.89	26.12	85.40
	± 1.46	± 3.61	± 1.55	± 4.92	± 1.50	± 3.50	± 1.73	± 6.28

pling, implying they can subsidize the acquisition of lower SNR contrasts. Therefore, a critical direction for future research is the dynamic optimization of these ratios within a fixed total scan time. For example, a protocol could allocate $R = 16\times$ to T_1 -MPRAGE, allowing for a more conservative $R = 9\times$ for FLAIR, ensuring signal integrity is preserved where it is most vulnerable without extending the total duration.

While this study establishes the technical feasibility of the proposed ssMRF framework using well-established MRI quantitative metrics (e.g., PSNR, SSIM), we recognize that these quantitative measures cannot fully capture the diagnostic confidence and accuracy required for clinical adoption. We acknowledge that radiologist interpretation remains the ultimate standard for evaluating synthetic imaging, and the absence of qualitative expert assessment is a limitation of the current work. Now that the synthesis fidelity has been validated in healthy volunteers, the essential next step is a formal reader study. Future work will focus on radiologist evaluation of subjective image quality, artifact perceptibility, and diagnostic accuracy in a diverse patient cohort with pathology to bridge the gap between technical viability and clinical practice.

Several promising avenues can be explored in future work to further enhance the capabilities and clinical utility of the proposed ssMRF framework. While this study demonstrates ssMRF in healthy volunteers, comprehensive validation in patient populations with pathology is an important next step. This is an active direction we are currently pursuing in collaboration with clinical partners. Expanding evaluation to a broader and more diverse cohort will enable assessment of ssMRF performance across a range of disease processes and patient characteristics, and is necessary to establish robustness and general applicability in clinical practice [19, 64]. Secondly, transitioning from slice-based synthesis to volumetric synthesis using advanced 3D or 2.5D architectures [65–67] and progressive volumetrization [68] techniques can significantly improve the model's capability in capturing and utilizing global contextual information and consistency across slices in the longitudinal dimension without sacrificing fine details. Furthermore, the learned deep model for MRF to contrast synthesis can be optimized during inference with subject-specific fine-tuning strategies via the proposed semi-supervised loss function, providing data-consistent synthesis and enhancing the accuracy of the resulting synthesized images [59]. These advancements could position ssMRF as a more versatile and potent tool in clinical settings, potentially improving its diagnostic capability.

The observations regarding the textural sharpness of Bloch synthesis versus the high-fidelity smoothness of the deep learning predictions suggest a potential avenue for hybrid modeling. While Bloch simulations are prone to noise and

contrast mismatches (particularly in FLAIR), they inherently preserve high-frequency structural information that data-driven models may smooth out. Future work could explore a hybrid synthesis framework where Bloch synthesis results are incorporated as additional physical priors or input channels to the network. Such an approach could synergistically combine the textural definition of physics-based simulations with the superior contrast accuracy and noise reduction capabilities of the semi-supervised ssMRF model.

Finally, we note that the subspace coefficient images used as input are derived from a simplified EPG-based dictionary that does not explicitly model certain physiological effects such as magnetization transfer (MT) or flow. Consequently, specific temporal signal components driven by these effects may lie partially outside the span of the basis. Although prior work [60] suggests that subspaces learned from idealized EPG simulations can sometimes remain effective for unmodeled deviations when the dominant temporal structure is strongly correlated with the true signal, this robustness is not guaranteed, particularly for effects like MT and flow. The data-driven synthesis network may still provide implicit compensation when these effects are reflected in the input features or present in the training data, but performance may vary across cases. Future iterations could reduce this mismatch by expanding the dictionary and subspace to explicitly incorporate MT and flow parameters.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s10334-026-01338-7>.

Funding This study is supported in part by GE Healthcare research funds and in part by NIH grants R01EB020613, R01MH116173, R01HD114719, and R01EB019437

Data availability Data and code to reproduce the above results will be shared in this repository upon acceptance: <https://github.com/Setsonomf/Lab/ssmrf>.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Ethical approval All in vivo data were acquired with informed consent as approved by the Institutional Review Board.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons

licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

References

- Ma D, Gulani V, Seiberlich N, Liu K, Sunshine JL, Duerk JL, Griswold MA (2013) Magnetic resonance fingerprinting. *Nature* 495(7440):187–192
- Panda A, Mehta BB, Coppo S, Jiang Y, Ma D, Seiberlich N, Griswold MA, Gulani V (2017) Magnetic resonance fingerprinting—an overview. *Curr Opin Biomed Eng* 3:56–66
- Weiskopf N, Edwards LJ, Helms G, Mohammadi S, Kirilina E (2021) Quantitative magnetic resonance imaging of brain anatomy and in vivo histology. *Nat Rev Phys* 3(8):570–588
- Bipin Mehta B, Coppo S, Frances McGivney D, Ian Hamilton J, Chen Y, Jiang Y, Ma D, Seiberlich N, Gulani V, Alan Griswold M (2019) Magnetic resonance fingerprinting: a technical review. *Magn Reson Med* 81(1):25–46
- Assländer J, Cloos MA, Knoll F, Sodickson DK, Hennig J, Lattanzi R (2018) Low rank alternating direction method of multipliers reconstruction for mr fingerprinting. *Magn Reson Med* 79(1):83–96
- Cauley SF, Setsompop K, Ma D, Jiang Y, Ye H, Adalsteinsson E, Griswold MA, Wald LL (2015) Fast group matching for mr fingerprinting reconstruction. *Magn Reson Med* 74(2):523–528
- Tropp J, Gilbert AC et al (2007) Signal recovery from partial information via orthogonal matching pursuit. *IEEE Trans Inf Theory* 53(12):4655–4666
- Hamilton JI, Jiang Y, Chen Y, Ma D, Lo W-C, Griswold M, Seiberlich N (2017) Mr fingerprinting for rapid quantification of myocardial t1, t2, and proton spin density. *Magn Reson Med* 77(4):1446–1458
- Ma D, Jiang Y, Chen Y, McGivney D, Mehta B, Gulani V, Griswold M (2018) Fast 3d magnetic resonance fingerprinting for a whole-brain coverage. *Magn Reson Med* 79(4):2190–2197
- Ye H, Cauley SF, Gagoski B, Bilgic B, Ma D, Jiang Y, Du YP, Griswold MA, Wald LL, Setsompop K (2017) Simultaneous multislice magnetic resonance fingerprinting (sms-mrf) with direct-spiral slice-grappa (ds-sg) reconstruction. *Magn Reson Med* 77(5):1966–1974
- Cao X, Liao C, Iyer SS, Wang Z, Zhou Z, Dai E, Liberman G, Dong Z, Gong T, He H, Zhong J, Bilgic B, Setsompop K (2022) Optimized multi-axis spiral projection MR fingerprinting with subspace reconstruction for rapid whole-brain high-isotropic-resolution quantitative imaging. *Magn Reson Med* 88(1):133–150. <https://doi.org/10.1002/mrm.29194>
- Liang Z-P (2007) Spatiotemporal imaging with partially separable functions. In: 2007 4th IEEE International Symposium on Biomedical Imaging: from Nano to Macro. IEEE, pp 988–991
- Zhao B, Setsompop K, Adalsteinsson E, Gagoski B, Ye H, Ma D, Jiang Y, Ellen Grant P, Griswold MA, Wald LL (2018) Improved magnetic resonance fingerprinting reconstruction with low-rank and subspace modeling. *Magn Reson Med* 79(2):933–942
- Seiler A, Nöth U, Hok P, Reiländer A, Maiworm M, Baudrexel S, Meuth S, Rosenow F, Steinmetz H, Wagner M et al (2021) Multiparametric quantitative mri in neurological diseases. *Front Neurol* 12:640239
- Filip P, Kokošová V, Valenta Z, Baláž M, Mangia S, Michaeli S, Vojtíšek L (2023) Utility of quantitative mri metrics in brain ageing research. *Front Aging Neurosci* 15:1099499
- Ma D, Jones SE, Deshmane A, Sakaie K, Pierre EY, Larvie M, McGivney D, Blümcke I, Krishnan B, Lowe M et al (2019) Development of high-resolution 3d mr fingerprinting for detection and characterization of epileptic lesions. *J Magn Reson Imaging* 49(5):1333–1346
- Gómez PA, Cencini M, Golbabae M, Schulte RF, Pirkel C, Horvath I, Fallo G, Peretti L, Tosetti M, Menze BH et al (2020) Rapid three-dimensional multiparametric mri with quantitative transient-state imaging. *Sci Rep* 10(1):13769
- Wang K, Doneva M, Meineke J, Amthor T, Karasan E, Tan F, Tamir JI, Yu SX, Lustig M (2023) High-fidelity direct contrast synthesis from magnetic resonance fingerprinting. *Magn Reson Med* 90(5):2116–2129
- Schauman S, Iyer SS, Yurt M, Cao X, Liao C, Zhong Z, Wang G, Zaharchuk G, Vasanaawala S, Setsompop K (2022) Toward a 1-minute high-resolution brain exam-mr fingerprinting with fast reconstruction and ml-synthesized contrasts 53
- Yurt M, Cao X, Zhou Z, Setsompop K, Vasanaawala S, Pauly J (2025) Foundation models for multimodal mri synthesis with language guidance. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 6759–6764
- Dar SU, Yurt M, Karacan L, Erdem A, Erdem E, Çukur T (2019) Image synthesis in multi-contrast mri with conditional generative adversarial networks. *IEEE Trans Med Imaging* 38(10):2375–2388. <https://doi.org/10.1109/TMI.2019.2901750>
- Tofts P (2005) Quantitative MRI of the brain: measuring changes caused by disease. Wiley, New York
- Hinshaw WS, Lent AH (1983) An introduction to nmr imaging: From the bloch equation to the imaging equation. *Proc IEEE* 71(3):338–350
- Weigel M (2015) Extended phase graphs: dephasing, rf pulses, and echoes—pure and simple. *J Magn Reson Imaging* 41(2):266–295
- Blystad I, Warntjes JBM, Smedby O, Landtblom A-M, Lundberg P, Larsson E-M (2012) Synthetic mri of the brain in a clinical setting. *Acta Radiol* 53(10):1158–1163
- Guenther C, Amthor T, Doneva M, Kozzerke S (2021) A unifying view on extended phase graphs and bloch simulations for quantitative mri. *Sci Rep* 11(1):21289
- Tanenbaum LN, Tsiouris AJ, Johnson AN, Naidich TP, DeLano MC, Melhem ER, Quarterman P, Parameswaran S, Shankaranarayanan A, Goyen M et al (2017) Synthetic mri for clinical neuroimaging: results of the magnetic resonance image compilation (magic) prospective, multicenter, multireader trial. *Am J Neuroradiol* 38(6):1103–1110
- Granberg T, Uppman M, Hashim F, Cananau C, Nordin L, Shams S, Berglund J, Forslin Y, Aspelin P, Fredrikson S et al (2016) Clinical feasibility of synthetic mri in multiple sclerosis: a diagnostic and volumetric validation study. *Am J Neuroradiol* 37(6):1023–1029
- Kleinloog JP, Mandija S, D'Agata F, Liu H, Heide O, Koktas B, Dankbaar JW, Keil VC, Vonken E-J, Jacobs SM et al (2023) Synthetic mri with magnetic resonance spin tomography in time-domain (mr-stat): results from a prospective cross-sectional clinical trial. *J Magn Reson Imaging* 57(5):1451–1461
- Ballester MAG, Zisserman AP, Brady M (2002) Estimation of the partial volume effect in mri. *Med Image Anal* 6(4):389–405
- Tang S, Fernandez-Granda C, Lannuzel S, Bernstein B, Lattanzi R, Cloos M, Knoll F, Assländer J (2018) Multicompartment magnetic resonance fingerprinting. *Inverse Prob* 34(9):094005
- Henkelman R, Stanisz G, Graham S (2001) Magnetization transfer in mri: a review. *NMR Biomed Int J Devoted Dev Appl Magn Reson Vivo* 14(2):57–64
- Qiu S, Ma S, Wang L, Chen Y, Fan Z, Moser FG, Maya M, Sati P, Sicotte NL, Christodoulou AG et al (2023) Direct synthesis of multi-contrast brain mr images from mr multitasking spatial factors using deep learning. *Magn Reson Med* 90(4):1672–1681
- Vargas M, Boto J, Delatre B (2016) Synthetic mr imaging sequence in daily clinical practice. *AJNR Am J Neuroradiol* 37(10):68

35. Ryu KH, Baek HJ, Moon JI, Choi BH, Park SE, Ha JY, Jeon KN, Bae K, Choi DS, Cho SB et al (2020) Initial clinical experience of synthetic mri as a routine neuroimaging protocol in daily practice: a single-center study. *J Neuroradiol* 47(2):151–160
36. Virtue P, Tamir JI, Doneva M, Yu SX, Lustig M (2018) Learning contrast synthesis from mr fingerprinting *Proc Intl Soc Mag Reson Med* 26:676
37. Nykänen O, Nevalainen M, Casula V, Isosalo A, Inkinen SI, Nikki M, Lattanzi R, Cloos MA, Nissi MJ, Nieminen MT (2023) Deep-learning-based contrast synthesis from mrf parameter maps in the knee joint. *J Magn Reson Imaging* 58(2):559–568
38. Cao T, Ma S, Wang N, Gharabaghi S, Xie Y, Fan Z, Hogg E, Wu C, Han F, Tagliati M et al (2022) Three-dimensional simultaneous brain mapping of t1, t2, and magnetic susceptibility with mr multitasking. *Magn Reson Med* 87(3):1375–1389
39. Zhu XJ (2005) Semi-supervised learning literature survey
40. Van Engelen JE, Hoos HH (2020) A survey on semi-supervised learning. *Mach Learn* 109(2):373–440
41. Kingma DP, Mohamed S, Jimenez Rezende D, Welling M (2014) Semi-supervised learning with deep generative models. *Adv Neural Inf Process Syst* 27
42. Grandvalet Y, Bengio Y (2004) Semi-supervised learning by entropy minimization. *Adv Neural Inf Process Syst* 17
43. Yurt M, Dalmaz O, Dar S, Ozbey M, Tinaz B, Oguz K, Çukur T (2022) Semi-supervised learning of mri synthesis without fully-sampled ground truths. *IEEE Trans Med Imaging* 41(12):3895–3906
44. Cheplygina V, De Bruijne M, Pluim JP (2019) Not-so-supervised: a survey of semi-supervised, multi-instance, and transfer learning in medical image analysis. *Med Image Anal* 54:280–296
45. Yaman B, Hosseini SAH, Moeller S, Ellermann J, Uğurbil K, Akçakaya M (2020) Self-supervised learning of physics-guided reconstruction neural networks without fully sampled reference data. *Magn Reson Med* 84(6):3172–3191
46. Liang P (2005) Semi-supervised learning for natural language. Ph.D. thesis, Massachusetts Institute of Technology
47. Chen L, Garcia F, Kumar V, Xie H, Lu J (2021) Industry scale semi-supervised learning for natural language understanding. [arXiv:2103.15871](https://arxiv.org/abs/2103.15871)
48. Dai AM, Le QV (2015) Semi-supervised sequence learning. *Adv Neural Inf Process Syst* 28
49. Turian J, Ratnoff L, Bengio Y (2010) Word representations: a simple and general method for semi-supervised learning. In: *Proceedings of the 48th annual meeting of the association for computational linguistics*, pp. 384–394
50. Zhai X, Oliver A, Kolesnikov A, Beyer L (2019) S4l: self-supervised semi-supervised learning. In: *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 1476–1485
51. Guillaumin M, Verbeek J, Schmid C (2010) Multimodal semi-supervised learning for image classification. In: *2010 IEEE computer society conference on computer vision and pattern recognition*. IEEE, pp 902–909
52. Honari S, Molchanov P, Tyree S, Vincent P, Pal C, Kautz J (2018) Improving landmark localization with semi-supervised learning. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1546–1555
53. Ouali Y, Hudelot C, Tami M (2020) Semi-supervised semantic segmentation with cross-consistency training. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp 12674–12684
54. Bai W, Oktay O, Sinclair M, Suzuki H, Rajchl M, Tarroni G, Glocker B, King A, Matthews PM, Rueckert D (2017) Semi-supervised learning for network-based cardiac mr image segmentation. In: *Medical image computing and computer-assisted intervention-MICCAI 2017: 20th international conference, Quebec City, QC, Canada, September 11–13, 2017, Proceedings, Part II* 20. Springer, pp 253–260
55. Madani A, Moradi M, Karagyris A, Syeda-Mahmood T (2018) Semi-supervised learning with generative adversarial networks for chest x-ray classification with ability of data domain adaptation. In: *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*. IEEE, pp 1038–1042
56. Luo X, Chen J, Song T, Wang G (2021) Semi-supervised medical image segmentation through dual-task consistency. In: *Proceedings of the AAAI conference on artificial intelligence*. pp 8801–8809
57. Cole EK, Ong F, Vasanaawala SS, Pauly JM (2021) Fast unsupervised mri reconstruction without fully-sampled ground truth data using generative adversarial networks. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp 3988–3997
58. Yurt M, Alkan C, Schauman S, Cao X, Liao C, Iyer S, Cukur T, Vasanaawala S, Pauly J, Setsompop K (2023) Semi-supervision for clinical contrast-weighted image synthesis from magnetic resonance fingerprinting. In: *Proc Int Soc Magn Reson Med*
59. Yurt M, Zhou Z, Liao C, Alkan C, Cao X, Wang N, Oscanoa J, Schauman S, Gao M, Li Z, Cukur T, Soares B, Syed A, Vasanaawala S, Pauly J, Setsompop K (2024) Unlocking data-consistent synthesis of clinical contrasts from magnetic resonance fingerprinting with semi-supervised learning. *Proc Int Soc Magn Reson Med*
60. Tamir JI, Uecker M, Chen W, Lai P, Alley MT, Vasanaawala SS, Lustig M (2017) T2 shuffling: sharp, multicontrast, volumetric fast spin-echo imaging. *Magn Reson Med* 77(1):180–195
61. Caruana R (1997) Multitask learning. *Mach Learn* 28:41–75
62. Zhang Y, Yang Q (2021) A survey on multi-task learning. *IEEE Trans Knowl Data Eng* 34(12):5586–5609
63. Uecker M, Tamir JI, Ong F, Lustig M (2016) The bart toolbox for computational magnetic resonance imaging. *Proc Intl Soc Magn Reson Med* 24:1
64. Hsieh JJ, Svalbe I (2020) Magnetic resonance fingerprinting: from evolution to clinical applications. *J Med Radiat Sci* 67(4):333–344
65. Kim S, Moon J, Chung H, Jang I (2024) Cyclic 2.5 d perceptual loss for cross-modal 3d image synthesis: T1 mri to tau-pet. [arXiv:2406.12632](https://arxiv.org/abs/2406.12632)
66. Hu Y, Kothapalli SV, Gan W, Sukstanskii AL, Wu GF, Goyal M, Yablonskiy DA, Kamilov US (2023) Diffgepci: 3d mri synthesis from mgre signals using 2.5 d diffusion model. [arXiv:2311.18073](https://arxiv.org/abs/2311.18073)
67. Qu T, Wang X, Fang C, Mao L, Li J, Li P, Qu J, Li X, Xue H, Yu Y et al (2022) M3net: a multi-scale multi-view framework for multi-phase pancreas segmentation based on cross-phase non-local attention. *Med Image Anal* 75:102232
68. Yurt M, Özbey M, Dar SU, Tinaz B, Oguz KK, Çukur T (2022) Progressively volumetrized deep generative models for data-efficient contextual learning of mr image recovery. *Med Image Anal* 78:102429