Article

Not peer-reviewed version

Meta-Learning–Driven Score-Based Generative Model for 0.5T-to-1.5T MRI Reconstruction

Yilin Su , Congcong Liu , Haifeng Wang , Yihang Zhou , Yuanyuan Liu , Jing Cheng , Qingyong Zhu , Qiegen Liu , Zhuoxu Cui * , Dong Liang *

Posted Date: 17 March 2026

doi: 10.20944/preprints202603.1356.v1

Keywords: low-field MRI; ill-posedness; meta-learning; unpaired learning; optimal-transport; scorebased diffusion model

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 17 March 2026 doi:10.20944/preprints202603.1356.v1

Article

Meta-Learning–Driven Score-Based Generative Model for 0.5T-to-1.5T MRI Reconstruction

Yilin Su [1,†], Congcong Liu [2,3,†], Haifeng Wang [2,4], Yihang Zhou [2], Yuanyuan Liu [2,4], Jing Cheng [2,4], Qingyong Zhu [2], Qiegen Liu [5], Zhuo-Xu Cui [2,*] and Dong Liang [1,2,4,6,*]

[1] School of Biomedical Engineering, Southern Medical University, Guangzhou, China
[2] Research Center for Medical AI, Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences, Shenzhen, China
[3] Shenzhen College of Advanced Technology, University of Chinese Academy of Sciences, Shenzhen, China
[4] Paul C. Lauterbur Research Center for Biomedical Imaging, Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences, Shenzhen, China
[5] School of Information Engineering, Nanchang University, Nanchang, China
[6] Pazhou Lab, Guangzhou, China
[*] Correspondence: zx.cui@siat.ac.cn (Z.-X.C.); dong.liang@siat.ac.cn (D.L.)
[†] These authors contributed equally to this work.

Abstract

Magnetic resonance imaging (MRI) acquired at low magnetic field strengths typically suffers from reduced signal-to-noise ratios (SNR), which leads to noticeable signal degradation compared with high-field MRI. As a result, reconstructing high-field-like images from low-field MRI data is a challenging task due to the inherently ill-posed nature of the problem. In addition, obtaining paired low-field and high-field MR images is often difficult in practical scenarios.To address these challenges, we propose a novel meta-learning framework with a two-stage mechanism. In the first stage, an optimal-transport-based meta-learner models the degradation process from high-field to low-field MRI and generates pseudo-paired datasets consisting of high-field and low-field images. In the second stage, a base learner solves the inverse problem of recovering high-field-like images from low-field MRI through an iterative regularization strategy, where the learned joint distribution of the pseudo-paired data serves as a prior.Experimental results demonstrate the capability of the proposed approach to generate 1.5T-like images from 0.5T MRI data. Both qualitative visualization and quantitative evaluations, conducted by comparing the reconstructed images with registered real 1.5T images, show that the proposed method produces images with SNR and contrast comparable to those of true 1.5T scans, even under a three-fold acceleration setting. Furthermore, the proposed method achieves superior performance compared with several mainstream approaches, including CycleGAN and Score-MRI.

Keywords: low-field MRI; ill-posedness; meta-learning; unpaired learning; optimal-transport; score-based diffusion model

---

1. Introduction

Magnetic resonance imaging (MRI) plays an essential role in modern clinical practice, as it provides high-resolution images of internal anatomical structures without requiring invasive procedures or exposure to ionizing radiation. However, conventional clinical MRI scanners typically operate at high magnetic field strengths and rely on copper coils that must be actively cooled using liquid helium. This requirement makes high-field MRI systems technically complex, expensive, and difficult to maintain. As a result, the accessibility of high-field MRI remains limited for many healthcare providers and patients worldwide, particularly in developing regions.In contrast, low-field MRI systems often employ permanent magnets, which are generally less costly and easier to operate and maintain. Consequently, low-field MRI has emerged as a promising and more accessible alternative to traditional high-field MRI scanners [1,2]. Despite these advantages, low-field MRI suffers from

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 17 March 2026 doi:10.20944/preprints202603.1356.v1

inherent limitations. In particular, the reduced magnetic field strength leads to a lower signal-to-noise ratio (SNR), which restricts image resolution and may hinder the visualization of fine anatomical details. Therefore, improving the image quality of low-field MRI is critical for facilitating its broader clinical adoption. In this work, we aim to address this limitation by reconstructing high-field-like (1.5T-like) MR images from low-field (0.5T) MRI data.

Nevertheless, achieving this objective presents several challenges. First, as the magnetic field strength decreases, both the signal-to-noise ratio and spatial resolution of MRI images tend to deteriorate. The transformation from high-field MRI to low-field MRI can therefore be regarded as a degradation process. Consequently, recovering high-field-like MRI images from low-field data constitutes a highly ill-posed inverse problem. Such problems are difficult to solve because they are typically characterized by non-uniqueness and instability [3–5]. Accurately learning the mapping from low-field MRI to high-field MRI is therefore a challenging task, making straightforward end-to-end deep learning approaches insufficient in many cases.

Another practical issue in MRI is the long acquisition time. Even in low-field MRI systems, prolonged scan durations remain a major limitation. As a result, considerable attention has been devoted to reconstructing high-quality images from undersampled k-space measurements in order to accelerate the acquisition process. However, in low-field MRI systems, the ill-posedness of the reconstruction problem becomes even more severe due to the combined effects of reduced SNR and k-space undersampling caused by accelerated imaging. These dual challenges are illustrated schematically in Figure 1.

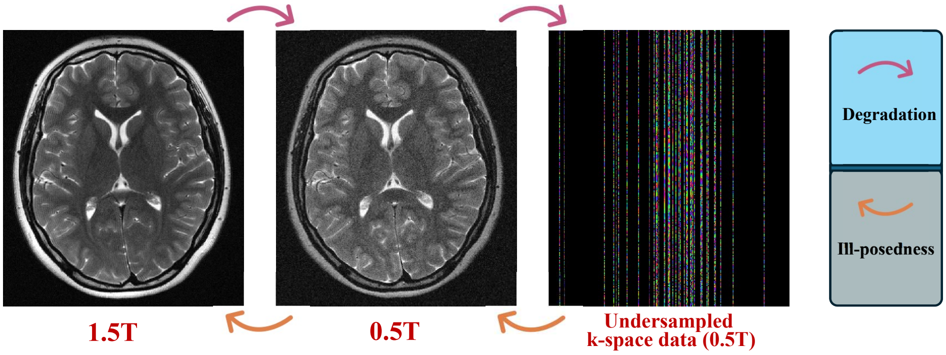

Figure 1. The diagram depicts the dual degradation of 1.5T MR images to 0.5T undersampled k-space data, as well as the double ill-posed nature of its reversal.

Furthermore, most existing approaches for reconstructing high-field-like MR images [6,7] rely on paired datasets containing both low-field and high-field MR images. Such paired data are often difficult to obtain in practical settings. Although unpaired learning strategies have been explored as an alternative, they typically focus on directly translating low-field MR images into high-field MR images. However, this strategy may not adequately address the underlying inverse problem involved in reconstructing high-field-like images from low-field data, as discussed above.

To tackle these challenges, it is necessary to develop a learning framework capable of effectively capturing the prior information shared between low-field and high-field MR images using unpaired datasets. Based on this learned prior, a regularization-based reconstruction model can then be constructed to mitigate the ill-posedness arising from reduced SNR, degraded resolution, and k-space undersampling.

1.1. Contributions

Based on the above challenges, this work makes the following contributions:

1. To tackle the ill-posed problem of reconstructing high-field-like MR images from low-field MRI in the absence of paired training data, we propose a novel meta-learning framework with a two-stage mechanism. In the first stage, an optimal-transport (OT)-based meta-learner models

the degradation process from high-field to low-field MRI and generates pseudo-paired high-field and low-field images. In the second stage, a base learner addresses the reconstruction inverse problem using an iterative regularization framework, where the learned joint distribution of the pseudo-paired images is incorporated as a regularizer.

2. Within this regularization framework, the proposed method further enables accelerated reconstruction, allowing high-field-like images to be recovered directly from undersampled low-field k-space measurements, which remains challenging for many existing approaches.
3. Experimental results on real 0.5T MRI data demonstrate that the proposed approach outperforms state-of-the-art unpaired learning methods in reconstructing high-field-like (1.5T-like) MR images, achieving improved contrast and higher SNR.

The remainder of this paper is organized as follows. Section 2 reviews related work. Section 3 introduces the proposed method. Implementation details are provided in Section 4. Experimental results on multiple datasets are presented in Section 5. Finally, Section 6 concludes the paper.

2. Related Work

2.1. Reconstruction of 7T-Like MR Images From 3T MRI

High-field MRI provides higher spatial resolution and improved tissue contrast in clinical practice, enabling more accurate and earlier diagnosis of brain diseases [8,9]. However, the high cost and technical complexity of high-field MRI systems significantly limit their accessibility, especially in developing regions. Therefore, generating high-quality images comparable to those acquired by high-field MRI from more affordable low-field MRI has become an important research direction.

Early studies addressed this problem from a super-resolution perspective, aiming to reconstruct high-resolution images from low-resolution inputs [10,11]. Nevertheless, the transformation from high-field MRI to low-field MRI involves not only resolution degradation but also reduced signal-to-noise ratio (SNR) and altered image contrast. To better model this relationship, [6] constructed a paired dictionary using 3T and 7T MR images to capture the correspondence between the two domains, enabling the reconstruction of 3T images with quality comparable to those acquired by 7T scanners.

However, reconstructing 1.5T-like images from 0.5T MRI is even more challenging. First, the SNR degradation is more severe: the SNR decreases by approximately 2.3 times from 7T to 3T, while it drops by nearly three times from 1.5T to 0.5T [12]. This leads to stronger image degradation and increased ill-posedness in the reconstruction problem. Second, the redundancy of acquisition coils in 0.5T systems is typically lower than that in 3T systems, further complicating the reconstruction process. In addition, existing approaches such as [6] rely on paired high-field and low-field datasets, which are difficult to obtain in practice. In contrast, this work considers the more realistic setting where paired data are unavailable and proposes an unpaired learning strategy to extract prior information from both 1.5T and 0.5T MR images. The learned priors are then incorporated as regularization to alleviate the ill-posedness of the reconstruction problem.

2.2. Unpaired Learning

In practice, acquiring paired medical images is often difficult due to various uncontrollable factors, such as patient motion and scanner variations. Consequently, learning correlation priors from unpaired data has become an active research topic. Existing approaches generally fall into two categories.The first category attempts to generate paired data based on physical models. For example, Bloch simulations can be used to synthesize low-field MR images from high-field MR images, after which conventional paired learning methods are applied. However, since such physical models rely on simplified assumptions, the simulated images may deviate from real measurements and introduce distortions [13]. The second category focuses on designing models that can be trained directly on unpaired data, such as optimal-transport-driven WGAN [14] and cycleGAN-based methods [15–17]. When applied to the reconstruction task considered in this work, these models attempt to learn a mapping between 0.5T and 1.5T MR images using unpaired datasets. However, it is important to note that the transformation from 1.5T to 0.5T MRI represents a degradation process. As will be shown later, a consistent bidirectional mapping between 0.5T and 1.5T MR images may not exist, which makes it difficult to directly apply such image translation methods to the reconstruction problem addressed in this paper.

2.3. Meta-Learning

Meta-learning is a powerful learning paradigm often described as "learning to learn" [18]. It is primarily applied in few-shot learning scenarios, where only a limited amount of labeled data is available. Traditional deep learning approaches typically learn a direct mapping between measurement data and labels, which becomes challenging when labeled samples are scarce. In contrast, meta-learning enables models to rapidly adapt and acquire new knowledge from a small number of examples [19]. Meta-learning has also been explored in medical imaging applications. For example, in [20], a meta-learning framework was introduced for low-dose CT reconstruction under limited labeled data. In that work, a teacher–student training paradigm was adopted: the teacher network, serving as a meta-learner, was trained using the small set of labeled data and then generated pseudo-labels for unlabeled low-dose CT images. The student network, acting as a base learner, subsequently learned to map low-dose CT images to the pseudo-labels produced by the teacher.

However, such mapping-based learning strategies are insufficient for the low-field MRI reconstruction problem considered in this work. Due to the severe ill-posedness introduced by reduced SNR and resolution, a reconstruction model requires stronger representation capability guided by data priors. Moreover, the semi-supervised learning strategy employed in the teacher–student framework is not applicable to the unpaired data setting addressed in this paper. Motivated by these observations, we develop a novel meta-learning framework specifically designed for the unpaired low-field MRI reconstruction scenario.

2.3.1. Score-Based Diffusion Models

Diffusion models have recently emerged as powerful generative frameworks capable of effectively capturing and exploiting data priors, and have achieved remarkable success across a wide range of applications [21–23]. In diffusion-based generative modeling, the original data are progressively perturbed by Gaussian noise through a forward diffusion process, which can be formulated using stochastic differential equations (SDEs) [24–27] or Markov chains [28]. As the perturbation proceeds, the data distribution gradually evolves toward pure noise. A score-matching strategy [29] is then employed to learn the score function associated with the intermediate distributions along the diffusion trajectory. Once the score function is learned, the original data can be recovered by reversing the diffusion process through reverse SDEs or reverse Markov chains.

The prior distribution learned by diffusion models can serve as an effective regularization in MRI reconstruction tasks. In particular, undersampled k-space measurements can be incorporated into the reverse diffusion process, leading to a conditional inverse SDE or Markov chain that iteratively solves the reconstruction problem. In this work, we leverage the capability of diffusion models to capture data priors and model the correlation between 1.5T and 0.5T MR images through a joint distribution. This learned joint distribution is incorporated as a regularization term in the reconstruction framework, enabling the recovery of high-quality 1.5T-like images from undersampled 0.5T MRI data.

Table 1. Mathematical symbols and their meanings.

Symbol	Meaning
$x^{1.5}$ or $x^{0.5}$	1.5T or 0.5T MR image
$y^{1.5}$ or $y^{0.5}$	1.5T or 0.5T undersampled k-space data
x^H	Hermitian transpose of x (conjugate and transpose)
$p(x)$	Probability density function (P.D.F.) of x
$p(x, y)$	Joint P.D.F. of x and y
δ_x	Dirac measure on x
α or β	Distributions that $x^{1.5}$ or $x^{0.5}$ follows, respectively
$\boldsymbol{T}$	Measurable map
$\boldsymbol{T}_\sharp\alpha$	The image measure of α through the map $\boldsymbol{T}$

3. Methodology

In this section, we first introduce a forward model that describes the transformation from 1.5T MR images to undersampled 0.5T k-space measurements. Based on this formulation, we then present a two-stage framework for reconstructing 1.5T-like images from undersampled 0.5T k-space data. In the first stage, the model learns from unpaired datasets to generate pseudo-paired samples. In the second stage, the joint distribution of the generated pairs is learned and incorporated into a regularization framework to reconstruct 1.5T MR images conditioned on undersampled 0.5T k-space observations.

3.1. Forward Model for the Degradation from 1.5T Images to 0.5T MR Data

The degradation process from 1.5T MR images to 0.5T MR images can be represented by the following forward model:

$$\mathbf{x}^{0.5} = \boldsymbol{f}(\mathbf{x}^{1.5}) \tag{1}$$

where $\mathbf{x}^{1.5}$ and $\mathbf{x}^{0.5}$ denote the high-field (1.5T) and low-field (0.5T) MR images, respectively, and $\boldsymbol{f}$ represents an unknown nonlinear degradation operator that maps high-field images to their low-field counterparts.

The acquisition process of 0.5T k-space measurements can be modeled as

$$\mathbf{y}^{0.5} = \boldsymbol{A}\mathbf{x}^{0.5} + \boldsymbol{n}, \tag{2}$$

where $\mathbf{y}^{0.5}$ denotes the undersampled k-space measurements at 0.5T, $\boldsymbol{A}$ represents the encoding operator, and $\boldsymbol{n}$ denotes Gaussian measurement noise. For multichannel MRI acquisition, the encoding operator can be expressed as $\boldsymbol{A} = \boldsymbol{PFS}$, where $\boldsymbol{S}$ denotes the coil sensitivity maps, $\boldsymbol{F}$ is the Fourier transform operator, and $\boldsymbol{P}$ represents the undersampling mask.

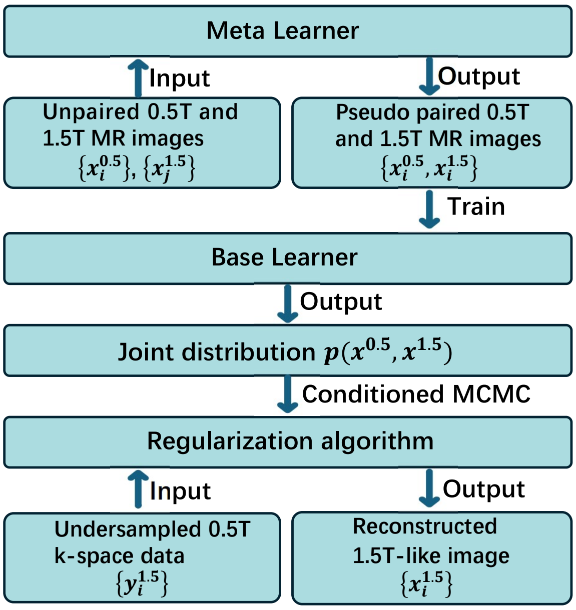

Figure 2. Illustration of the proposed meta-learning framework for reconstructing 1.5T-like MR images from 0.5T MRI. (a) In the first stage, a meta-learner models the degradation process from high-field to low-field MRI using unpaired datasets $\mathbf{x}_i^{0.5}$ and $\mathbf{x}_j^{1.5}$, and generates pseudo-paired samples $\mathbf{x}_i^{0.5}, \mathbf{x}_i^{1.5}$. (b) In the second stage, a base learner reconstructs a 1.5T-like MR image from undersampled 0.5T k-space measurements $\mathbf{y}_i^{0.5}$ within an iterative regularization framework, where the learned joint distribution $p(\mathbf{x}_i^{0.5}, \mathbf{x}_i^{1.5})$ is incorporated as a regularizer.

3.2. OT Driven Pseudo-Paired Data Generator

To address the inverse problems described in (1) and (2), conventional end-to-end deep learning approaches typically attempt to approximate the mapping from $\mathbf{x}^{0.5}$ to $\mathbf{x}^{1.5}$ using neural networks. However, in scenarios where paired training data are unavailable, such as the setting considered in this work, GAN-based methods can be adopted to relax the requirement for paired samples. Instead of minimizing a pointwise loss between paired images, GAN-based frameworks elevate the objective to matching distributions of data.

Let $\mathbf{x}^{1.5}$ and $\mathbf{x}^{0.5}$ denote samples drawn from distributions α and β, respectively. The goal of a GAN-based approach is then to learn an optimal transport operator $\boldsymbol{T}^{\beta\to\alpha}$ that maps the distribution β to α, satisfying

$$\boldsymbol{T}_{\sharp}^{\beta\to\alpha}(\beta) = \alpha. \tag{3}$$

Under this formulation, the operator $\boldsymbol{T}^{\beta\to\alpha}$ enables the generation of high-field images from low-field images, i.e., $\boldsymbol{T}^{\beta\to\alpha}(\mathbf{x}^{0.5}) = \mathbf{x}^{1.5}$.

However, due to the ill-posed nature of the degradation operator f, accurately learning the mapping from $\mathbf{x}^{0.5}$ to $\mathbf{x}^{1.5}$ is extremely challenging in practice, as will be confirmed in the experimental results. In contrast, our empirical observations indicate that the degradation process from $\mathbf{x}^{1.5}$ to $\mathbf{x}^{0.5}$, namely f, is considerably easier to approximate. Motivated by this observation, we propose a two-stage framework for solving the inverse problems (1) and (2). Specifically, the first stage learns the degradation mapping f to generate pseudo-paired samples, while the second stage learns the joint distribution of these pseudo pairs and incorporates it as a regularization prior for solving the inverse problem.

For computational tractability, the optimal transport from α to β is typically obtained by solving the Kantorovich dual formulation [30]:

$$\boldsymbol{T}^{\alpha\to\beta} = \arg\min_{\boldsymbol{T}} \max_{\varphi\in\mathrm{Lip}_1} \int \varphi(\boldsymbol{T}(\mathbf{z}))d\beta(\mathbf{z}) - \int \varphi(\mathbf{x})d\alpha(\mathbf{x}), \tag{4}$$

where Lip_1 denotes the set of functions that are uniformly Lipschitz continuous with Lipschitz constant 1. In practice, this constraint is often enforced using spectral normalization [31]. Numerically, the optimal transport operator $\boldsymbol{T}^{\alpha\to\beta}$ and the corresponding maximizer φ^* can be obtained through an alternating minimization–maximization optimization procedure. Once the optimal transport $\boldsymbol{T}^{\alpha\to\beta}$ is obtained, pseudo-paired samples of 1.5T and 0.5T MR images can be generated using the following procedure.

3.3. Score-Based Reconstruction Model

After obtaining an accurate approximation of $\boldsymbol{f}$, a common strategy is to combine the nonlinear inverse problems (1) and (2) into the following observation model:

$$\mathbf{y}^{0.5} = \boldsymbol{A}\boldsymbol{f}(\mathbf{x}^{1.5}) + \boldsymbol{n} \tag{5}$$

From a statistical perspective, solving the above inverse problem can be formulated as posterior sampling, which requires evaluating the gradient of the following posterior distribution:

$$\log p(\tilde{\mathbf{x}}^{1.5}|\mathbf{y}^{0.5}) = \log p(\mathbf{y}^{0.5}|\tilde{\mathbf{x}}^{1.5}) + \log p(\tilde{\mathbf{x}}^{1.5}) + C \tag{6}$$

To estimate the regularization term $\log p(\tilde{\mathbf{x}}^{1.5})$, diffusion models perturb a clean 1.5T image $\mathbf{x}^{1.5}$ with Gaussian noise at multiple scales and learn $\nabla \log p(\tilde{\mathbf{x}}^{1.5})$ along the perturbation trajectory using the score-matching method, where $\tilde{\mathbf{x}}^{1.5}$ denotes the perturbed image. However, due to the nonlinearity of $\boldsymbol{f}$ (i.e., $\boldsymbol{T}^{\alpha\to\beta}$), it is difficult to characterize how the perturbation noise in $\tilde{\mathbf{x}}^{1.5}$ propagates through $\boldsymbol{f}$. Consequently, evaluating $\log p(\mathbf{y}^{0.5}|\tilde{\mathbf{x}}^{1.5})$ becomes challenging. This issue motivates us to develop a new approach for reconstructing $\mathbf{x}^{1.5}$ from $\mathbf{y}^{0.5}$.

To avoid computing $\log p(\mathbf{y}^{0.5}|\tilde{\mathbf{x}}_t^{1.5})$, we instead consider the following posterior formulation to solve the inverse problems (1) and (2) without explicitly combining them:

$$\log p(\tilde{\mathbf{x}}^{1.5}, \tilde{\mathbf{x}}^{0.5}|\mathbf{y}^{0.5}) = \log p(\mathbf{y}^{0.5}|\tilde{\mathbf{x}}^{0.5}) + \log p(\tilde{\mathbf{x}}^{0.5}, \tilde{\mathbf{x}}^{1.5}) + C \tag{7}$$

According to the forward model (2), the data consistency term is given by

$$\nabla \log p(\mathbf{y}^{0.5}|\tilde{\mathbf{x}}^{0.5}) = -\frac{\boldsymbol{A}^{\boldsymbol{H}}(\boldsymbol{A}\mathbf{x}_t - \mathbf{y})}{\gamma^2 + \varepsilon^2}, \tag{8}$$

where γ denotes the noise level of $\mathbf{n}$ in (2), and ε represents the scale of the perturbation noise in $\tilde{\mathbf{x}}$. Using the pseudo-paired data generated by Algorithm 1, the joint distribution $p(\tilde{\mathbf{x}}^{1.5}, \tilde{\mathbf{x}}^{0.5})$ can be approximated via the score-matching method.

Algorithm 1 OT Driven Pseudo-Paired Data Generator.

```
Input: $T^{\alpha\to\beta}$ and 1.5T MR images $\{\mathbf{x}_i^{1.5}\}_{i=1}^N$;
for $i = 1, 2, \ldots, N$ do
    $\mathbf{x}_i^{0.5} = T^{\alpha\to\beta}(\mathbf{x}_i^{1.5})$
end for
Output: $\{(\mathbf{x}_i^{0.5}, \mathbf{x}_i^{1.5})\}_{i=1}^N$.
```

More specifically, we perturb $X_0 := [\mathbf{x}^{0.5}, \mathbf{x}^{1.5}]$ using Gaussian noise with scales $\varepsilon_{i}{}_{i=1}^{T}$ satisfying $\varepsilon_1 < \varepsilon_2 < \cdots < \varepsilon_T$. Let $p(\widetilde{X}_i, |X_0) = \mathcal{N}(\widetilde{X}_i|X_0, \varepsilon_i^2\boldsymbol{I})$, and the corresponding perturbed data distribution is given by $p(\widetilde{X}i) = \int p(\widetilde{X}i|X_0)\mathrm{d}\pi(X_0)$, where $\pi := (\boldsymbol{T}^{\alpha\to\beta}, \boldsymbol{I})_\sharp\alpha$. If ε_1. If ε_1 is chosen sufficiently small, then $p_{\varepsilon_1}(X) = p(X_0)$ approximately holds.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 17 March 2026 doi:10.20944/preprints202603.1356.v1

Based on the score-matching principle, we estimate $\nabla \log p(\widetilde{X}i)$ at different noise scales by training a joint score function $\boldsymbol{s_\phi}(\widetilde{X}i, \varepsilon_i)$ with the following objective:

$$\boldsymbol{\phi}^* = \arg\min_{\boldsymbol{\phi}} \frac{1}{2L} \sum_{i=1}^{L} \mathbb{E}_{\pi(X_0)} \mathbb{E}_{p(\widetilde{X}_i|X_0)} \left[\left\| \varepsilon_i \boldsymbol{s_\phi}(\widetilde{X}_i, \varepsilon_i) + \frac{\mathbf{z}}{\varepsilon_i} \right\|^2 \right] \tag{9}$$

where $\mathbf{z}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})$.

Once we have trained a scoring function, we can carry out conditionally Langevin MCMC sampling based on equation (7), that is

$$\begin{aligned}
\widetilde{X}_{i+1} &= \widetilde{X}_i + \frac{\eta_i}{2} \nabla \log p(\widetilde{X}_i | \mathbf{y}^{0.5}) + \sqrt{\eta_i} \mathbf{z}_i \\
&= \widetilde{X}_i + \frac{\eta_i}{2} (\nabla \log p(\widetilde{X}_i) + \nabla \log p(\mathbf{y}^{0.5} | \widetilde{\mathbf{x}}_i^{0.5})) + \sqrt{\eta_i} \mathbf{z}_i \\
&= \widetilde{X}_i + \frac{\eta_i}{2} \left(\boldsymbol{s_\phi}(\widetilde{X}_i, \varepsilon_i) - \frac{\boldsymbol{A^H}(\boldsymbol{A}\widetilde{\mathbf{x}}_i^{0.5} - \mathbf{y}^{0.5})}{\gamma^2 + \varepsilon_i^2} \right) + \sqrt{\eta_i} \mathbf{z}_i
\end{aligned} \tag{10}$$

In particular, the choice of parameter η_i follows that of literature [32], and the conditional Langevin MCMC sampling is detailed in Algorithm 2.

Algorithm 2 Conditional Langevin MCMC Sampling.

```
Input: $\mathbf{y}^{0.5}$, $\boldsymbol{\phi}^*$, $\{\varepsilon_i\}_{i=1}^{L}$, $\epsilon$ and $T$;
Initialize: $X_0 := [\widetilde{\mathbf{x}}_0^{0.5}, \widetilde{\mathbf{x}}_0^{1.5}]$;
for $i = L, L-1, \ldots, 1$ do
    $\eta_i = \epsilon \cdot \varepsilon_i^2 / \varepsilon_L^2$
    for $t = 0, 2, \ldots, T-1$ do
        Draw $\mathbf{z}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})$:
            $\widetilde{X}_{t+1} = \widetilde{X}_t + \frac{\eta_i}{2}\left(\boldsymbol{s}_{\boldsymbol{\phi}^*}(\widetilde{X}_t, \varepsilon_i) - \frac{\boldsymbol{A^H}(\boldsymbol{A}\widetilde{\mathbf{x}}_t^{0.5} - \mathbf{y}^{0.5})}{\gamma^2 + \varepsilon_i^2}\right) + \sqrt{\eta_i}\mathbf{z}_t$
    end for
    $\widetilde{X}_0 = \widetilde{X}_T$
end for
Output: $\widetilde{X}_T = [\widetilde{\mathbf{x}}_T^{0.5}, \widetilde{\mathbf{x}}_T^{1.5}]$.
```

4. Implementation

The evaluation was conducted using MRI data acquired from SuperMark 1.5T and SuperMark 0.5T scanners.

4.1. Data Acquisition

4.1.1. Anke 0.5T MRI Data

Brain raw data were acquired using a 0.5T Anke MRI system (SuperMark 0.5T) equipped with a 4-channel receive coil. Both T1-weighted and T2-weighted images were obtained using a conventional Cartesian 2D fast spin echo (FSE) sequence. The imaging parameters were as follows: TR/TE = 665/19 ms for T1-weighted imaging and TR/TE = 5350/112 ms for T2-weighted imaging. The field of view (FOV) was 230 × 230 mm^2 with a slice thickness of 5 mm and a single signal average. The acquisition matrix size was 320 × 320. The echo train length (ETL) was 3 for T1-weighted scans and 13 for T2-weighted scans. The in-plane spatial resolution was 0.71 × 0.71 mm^2, and the receiver bandwidth was 130 Hz/Pixel.

4.1.2. Anke 1.5T MRI Data

Brain raw data were also collected using a 1.5T Anke MRI scanner (SuperMark 1.5T) equipped with an 8-channel receive coil. Similar to the 0.5T acquisition, T1-weighted and T2-weighted images were obtained using a Cartesian 2D fast spin echo (FSE) sequence. The imaging parameters were set as follows: TR/TE = 808/10 ms for T1-weighted imaging and TR/TE = 3700/94 ms for T2-weighted imaging. The field of view (FOV) was 230×230 mm^2 with a slice thickness of 5 mm and a single signal average. The matrix size was 320×320. The echo train length (ETL) was 3 for T1-weighted scans and 19 for T2-weighted scans. The in-plane resolution was 0.71×0.71 mm^2, and the receiver bandwidth was 240 Hz/Pixel.

4.1.3. Training, Testing, and Validation Data Sets

MRI scans were performed on 47 subjects using both 1.5T and 0.5T MRI systems, and T1-weighted as well as T2-weighted images were acquired for each subject. In total, 47 datasets were obtained, each containing four image types: (0.5T-T1W, 0.5T-T2W, 1.5T-T1W, and 1.5T-T2W). Each dataset includes 19 axial slices.Although every subject was scanned using both field strengths, the images from the 0.5T and 1.5T scanners could not be perfectly aligned due to differences in scanner positioning, slice selection, and other acquisition conditions between the two systems. As a result, the images from the same subject at different field strengths cannot be regarded as strictly paired data.To enable subsequent quantitative evaluation through image registration, we ensured relatively consistent scanner positioning and slice selection for 7 subjects scanned at both 0.5T and 1.5T. Among these 7 datasets, 5 were randomly selected as the test set, while the remaining 2 were used as the validation set. The data from the other 40 subjects were used for training, where strict consistency in scanner positioning and slice selection was not required.It should be noted that training separate models for different imaging modalities may lead to overfitting due to the limited size of the dataset. Therefore, the training data consist of a mixture of both T1-weighted and T2-weighted images.

4.2. Network Architecture and Training

We first train the pseudo-paired data generator $\boldsymbol{T}^{\alpha \to \beta}$ using the optimal transport (OT) framework. As reported in [17], constructing bidirectional mappings between the distributions β and α during network training can impose mutual constraints and thus act as a form of regularization. Based on this idea, we introduce $\boldsymbol{T}^{\beta \to \alpha}$ into the OT model to further enforce the accuracy of $\boldsymbol{T}^{\alpha \to \beta}$. The discriminator and generator loss functions, denoted as $LossD$ and $LossG$, were optimized using the ADAM optimizer with parameters $\beta_1 = 0.5$ and $\beta_2 = 0.999$. The mini-batch size was set to 1. Training was performed with a learning rate of 10^{-3}, and an ExponentialLR scheduler with $\gamma = 0.95$ was adopted.

For the score-based reconstruction model, we employed the NCSNv2 network [32] to learn the score function. The parameters were configured as $\varepsilon_L = 50$, $\varepsilon_1 = 0.0066$, and the number of classes was set to 266. Exponential Moving Average (EMA) was enabled with a decay rate of 0.999. The loss function in (9) was optimized using the ADAM optimizer with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. The mini-batch size was set to 1, and the model was trained with a learning rate of 10^{-4}.

For the CycleGAN model, the maximum number of training epochs was set to 200. During validation, training was stopped once the pseudo 0.5T images generated by G_{hl} visually approximated real 0.5T images and the quantitative metric reached a predefined threshold (average PSNR > 20). In practice, convergence was achieved at the 50th epoch. For the score-based reconstruction model, the maximum training epoch was set to 500. Training was terminated when the reconstructed pseudo 1.5T images on the validation set visually resembled real 1.5T images and the quantitative metric satisfied the threshold (average PSNR > 20). The training process ultimately stopped at the 350th epoch.

All models were implemented on an Ubuntu 20.04 operating system using the open-source PyTorch 1.10 framework [33] with CUDA 11.3 and CUDNN support, and trained on an NVIDIA A100 Tensor Core GPU with 80 GB memory. The numbers of parameters for the corresponding networks are summarized in Table 2, and the overall training procedures are illustrated in Figure 3.

Table 2. Parameter count of different networks

Network	Parameters
ResNet	11371906
Pyramid Net	2763713
NCSNv2	29696388

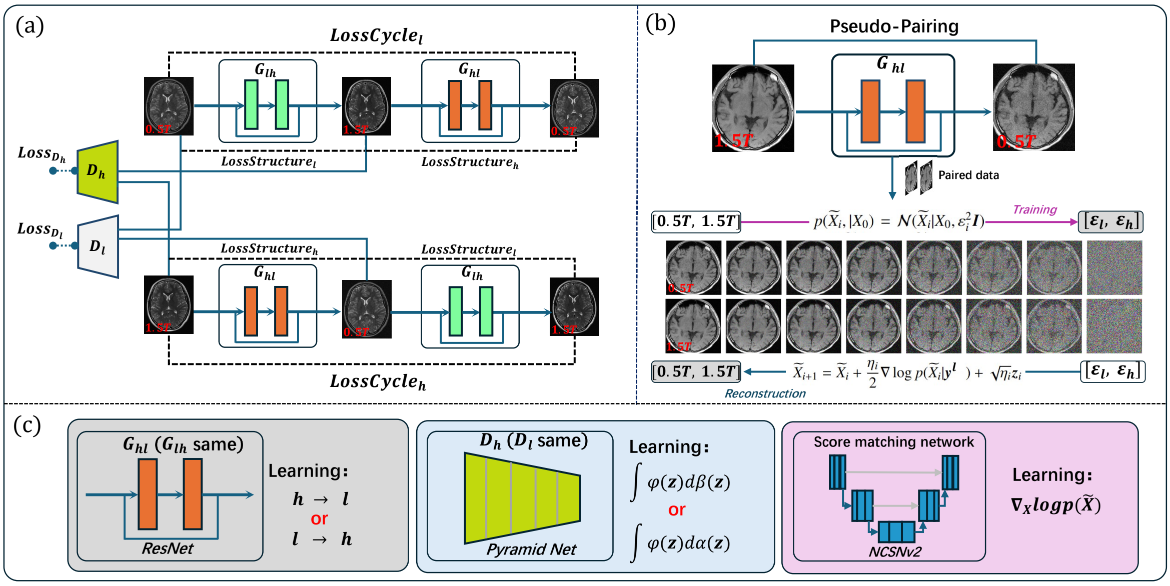

$p(\widetilde{X}_i|X_0) = \mathcal{N}(\widetilde{X}_i|X_0, \varepsilon_i^2 I)$

$\widetilde{X}_{i+1} = \widetilde{X}_i + \frac{\eta_i}{2}\nabla \log p(\widetilde{X}_i|y^l) + \sqrt{\eta_i} z_i$

$\int \varphi(z)d\beta(z)$ $\int \varphi(z)d\alpha(z)$

$\nabla_X \log p(\widetilde{X})$

Figure 3. Illustration of the training procedures for (a) the OT-based pseudo-paired $(\mathbf{x}^{0.5}, \mathbf{x}^{1.5})$ MR image generator and (b) the score-based reconstruction model that learns the joint distribution $p(\mathbf{x}^{0.5}, \mathbf{x}^{1.5})$.

4.3. Experimental Setup

First, we compare the proposed approach with the mapping-based method CycleGAN [15] in order to verify our claim that such mapping-learning strategies are not well suited for the scenario considered in this work. In addition, we approximate $p(\mathbf{y}^{0.5}|\tilde{\mathbf{x}}^{1.5})$ in model (6) using a Gaussian distribution, which requires modeling f as a noisy transformation process. This approximation, referred to as Score-MRI, serves as a baseline to highlight the advantages of the proposed model in (7).

In the experiments, we first investigate the ill-posedness arising from undersampling in (2) through an ablation study. Under this setting, the task reduces to generating fully sampled 1.5T MR images from fully sampled 0.5T MR images. Subsequently, we evaluate the effect of dual ill-posedness by introducing 3-fold undersampling to the 0.5T data. Considering that the 0.5T acquisition employs only 4 channels, a 3-fold undersampling rate is already close to the practical limit. This experiment enables us to assess the performance of different methods under more challenging reconstruction conditions.

4.4. Performance Evaluation

In this study, the quality of reconstructed 1.5T-like MR images from 0.5T data is primarily assessed through visual inspection. Furthermore, because the test set maintains relatively consistent slice positions between the 0.5T and 1.5T scans, the 1.5T images are additionally registered to the corresponding 0.5T images to serve as reference labels for quantitative evaluation. Since our goal is to determine whether the reconstructed 1.5T-like images can approximate real 1.5T images in terms of signal-to-noise ratio and image contrast, the peak signal-to-noise ratio (PSNR) is adopted as the quantitative evaluation metric.

5. Results

In this section, we evaluate the performance of the proposed Meta-Score-MRI method and validate the aforementioned claims through a series of experimental studies.

5.1. Reconstruction of 1.5T-like MR Image from Full-Sampled 0.5T MRI Data

This section evaluates several methods for reconstructing 1.5T-like MR images from fully sampled 0.5T MRI data, aiming to address the ill-posedness described in (2). The reconstruction results for 1.5T-like T2-weighted images are presented in Figure 4. It can be observed that the 0.5T images exhibit noticeably lower SNR and reduced contrast compared with the corresponding 1.5T images. Among the compared methods, CycleGAN generated images with contrast levels comparable to real 1.5T images; however, visible artifacts appear in the regions indicated by the red arrows, suggesting that a reliable OT mapping from 0.5T images to 1.5T images does not exist. The Score-MRI approach improves the SNR of the reconstructed images but fails to capture the contrast variations, which is consistent with its modeling assumption. In comparison, the proposed Meta-Score-MRI method produces reconstructed images that closely match the real 1.5T images, with minimal discrepancies in both SNR and contrast.

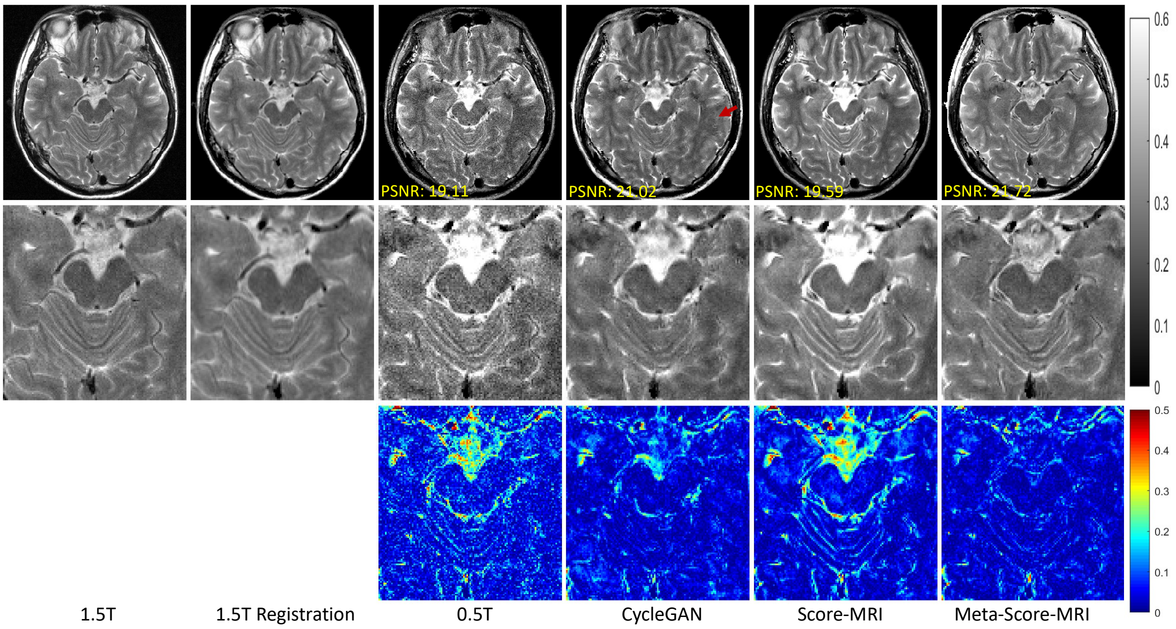

Figure 4. Reconstruction of 1.5T-like T2-weighted MR images from fully sampled 0.5T T2-weighted MRI data is shown. The first row presents the original 1.5T images, while the second row shows the corresponding 1.5T images after registration with the 0.5T images. The third row displays the acquired 0.5T MR images. The reconstructed images are presented in rows four through six. A zoomed-in view is provided in the second row, and the third row also includes error maps representing the differences between the reconstructed images from 0.5T data and the registered 1.5T images. PSNR values for each slice are indicated at the corners. The grayscale for the reconstructed images and the color bar for the error maps are shown on the right side of the figure.

The reconstruction results of 1.5T T1-weighted images are presented in Figure 5. The images generated by CycleGAN show only limited improvement in SNR and contrast. Score-MRI improves the SNR but largely preserves the contrast characteristics of the 0.5T images. In comparison, the proposed method produces contrast that is most consistent with the 1.5T reference and achieves visually comparable SNR. Furthermore, the zoomed-in view shows that Score-MRI yields blurred structural details, whereas the proposed method is able to preserve these details more accurately.

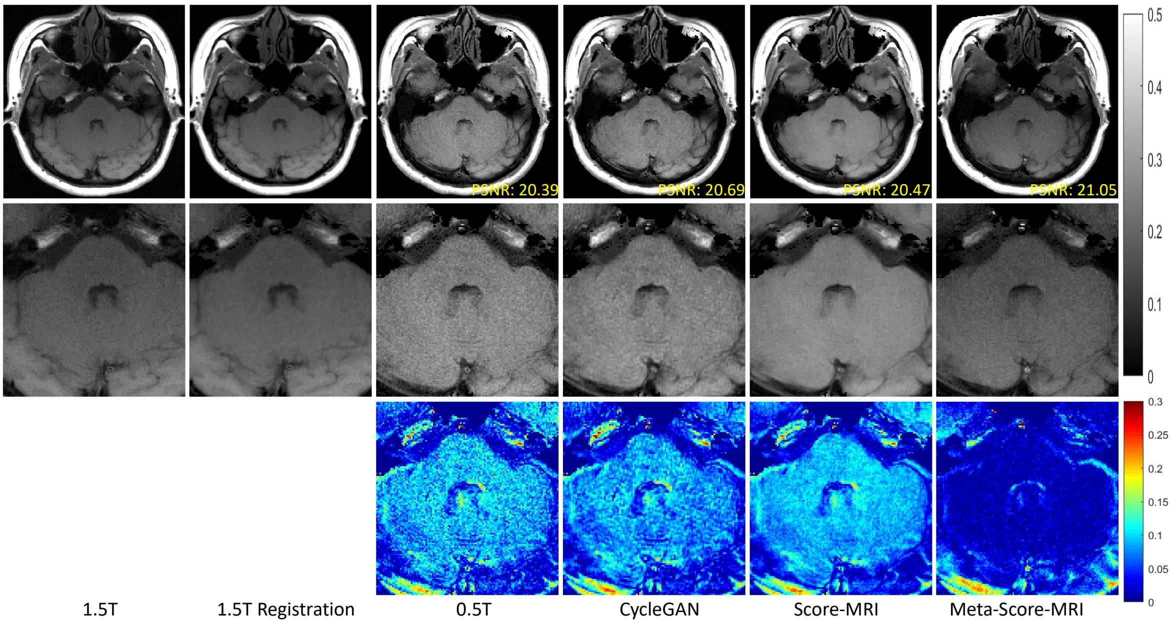

Figure 5. Reconstruction of 1.5T-like T1-weighted MR images from fully sampled 0.5T T1-weighted MRI data. The first row shows the original 1.5T MR images, while the second row presents the corresponding 1.5T images after registration to the 0.5T images. The third row displays the acquired 0.5T MR images. The reconstructed results are illustrated in rows four to six. A zoomed-in view is provided in the second row, and the third row also shows the error maps between the reconstructed images from 0.5T data and the registered 1.5T images. The PSNR value of each slice is indicated in the corner of the images. The grayscale for the reconstructed images and the color bar for the error maps are shown on the right side of the figure.

To further evaluate the effectiveness of the proposed method, we quantified the differences in contrast and SNR between the reconstructed 1.5T-like images produced by different methods and the registered real 1.5T images using PSNR. The quantitative comparisons of the competing methods are summarized in Table 3. These results are consistent with the observations from the visual analysis. The proposed method achieves the best quantitative performance for both T1- and T2-weighted images. Notably, in the reconstruction of T2-weighted images, the proposed method outperforms the second-best method by approximately 1 dB. In contrast, the quantitative improvement for T1-weighted image reconstruction is less pronounced. This can be attributed to the characteristics of the collected dataset, where the contrast variation of T2-weighted images across different field strengths is more substantial, whereas the contrast variation of T1-weighted images is relatively limited.

Overall, both the visual comparisons and quantitative evaluations consistently demonstrate the effectiveness of the proposed method for reconstructing 1.5T-like MR images from fully sampled 0.5T MRI data.

Table 3. Quantitative comparison of different methods for reconstructing 1.5T-like MR images from fully sampled 0.5T MRI data. The table reports the mean and standard deviation of the PSNR values for the reconstructed images produced by each method on the test set. The best results are highlighted in red.

Datasets & Methods		Quantitative Evaluation PSNR
T2W Recon	0.5T	19.80±0.85
	CycleGAN	20.63±0.96
	Score-MRI	20.10±1.02
	Meta-Score-MRI	21.54±1.20
T1W Recon	0.5T	21.99±1.51
	CycleGAN	22.49±1.79
	Score-MRI	22.17±1.58
	Meta-Score-MRI	22.52±1.56

5.2. *Reconstruction of 1.5T-like MR Image from Undersampled 0.5T MRI Data*

In this section, we reconstruct 1.5T-like MR images from 3-fold undersampled 0.5T k-space data and evaluate different methods for addressing the dual ill-posed problems described in Equations (1) and (2). Figure 6 presents the reconstruction results from undersampled 0.5T T2-weighted data. The zero-filled 0.5T image contains substantial noise and pronounced aliasing artifacts, making this task more challenging than the previous experiments. The images reconstructed by CycleGAN still exhibit noticeable aliasing artifacts that obscure anatomical details. In contrast, both Score-MRI and the proposed method effectively suppress these artifacts. However, as indicated by the red arrows, Score-MRI introduces spurious structures and retains the image characteristics of 0.5T MRI. By comparison, the proposed method accurately recovers fine details and produces images with SNR and contrast that closely match those of the original 1.5T images.

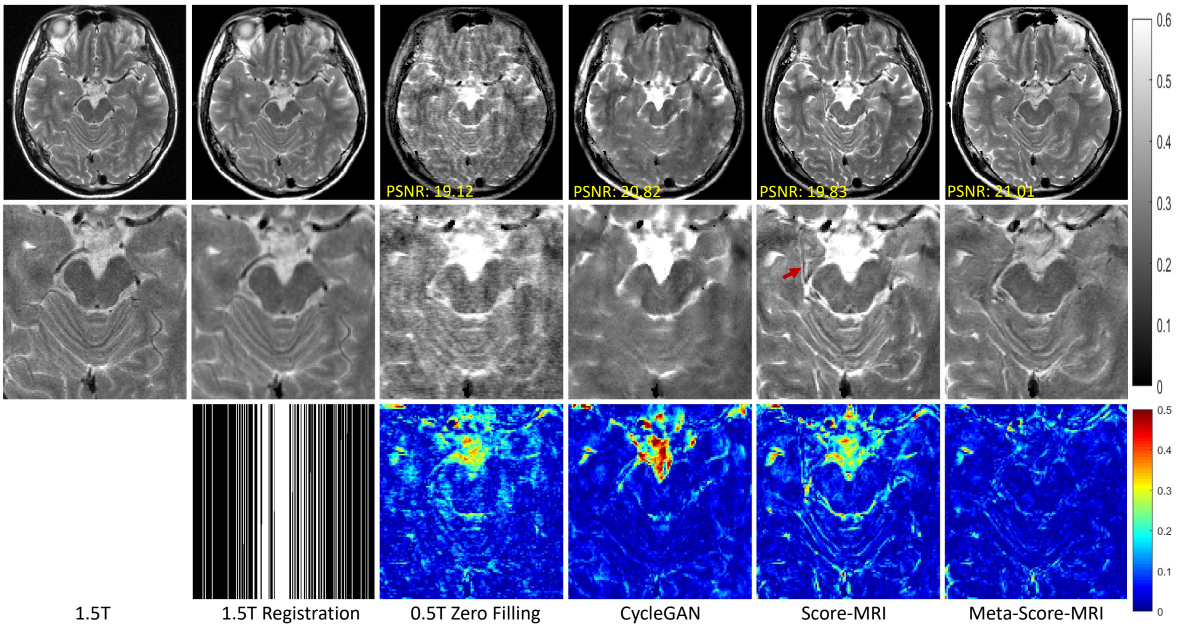

Figure 6. Reconstruction of 1.5T-like T2-weighted MR images from undersampled 0.5T T2-weighted k-space data with an acceleration factor of $R = 3$. The first row shows the original 1.5T MR images, while the second row presents the corresponding 1.5T images after registration with the 0.5T images. The third row displays the zero-filled 0.5T MR images. The reconstructed results are shown in rows four to six. A zoomed-in view is provided in the second row, and the third row also includes error maps between the reconstructed images from 0.5T data and the registered 1.5T images. The PSNR values for each slice are indicated in the image corners. The grayscale for the reconstructed images and the color bar for the error maps are shown on the right side of the figure.

The reconstruction results from 3-fold undersampled 0.5T T1-weighted k-space data are presented in Figure 7. The images reconstructed by CycleGAN show pronounced aliasing artifacts that obscure anatomical details. Although Score-MRI alleviates these artifacts, the contrast of the reconstructed images largely remains at the 0.5T level, resulting in blurred structures. In comparison, the proposed method effectively suppresses aliasing artifacts and produces images with SNR, contrast, and structural details comparable to those of the 1.5T reference images. The quantitative PSNR results summarized in Table 4 further demonstrate that the proposed method achieves superior performance compared with the competing approaches. These results highlight the effectiveness of our method in addressing the dual ill-posedness caused by low field strength and k-space undersampling.

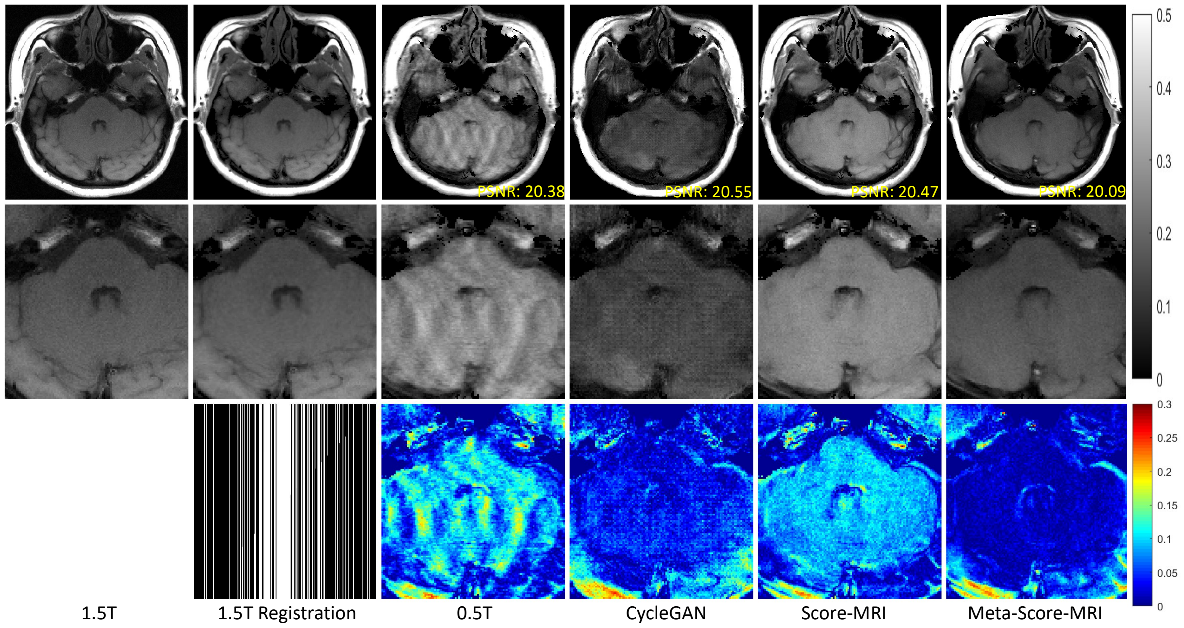

Figure 7. Reconstruction of 1.5T-like T1-weighted MR images from undersampled 0.5T T1-weighted k-space data with an acceleration factor of $R = 3$. The first row shows the original 1.5T MR images, while the second row presents the corresponding 1.5T images after registration with the 0.5T images. The third row displays the zero-filled 0.5T MR images. The reconstructed results are shown in rows four to six. A zoomed-in view is provided in the second row, and the third row also includes error maps between the reconstructed images from 0.5T data and the registered 1.5T images. The PSNR values for each slice are indicated in the image corners. The grayscale for the reconstructed images and the color bar for the error maps are shown on the right side of the figure.

Table 4. Quantitative comparison of different methods for reconstructing 1.5T-like MR images from 3-fold undersampled 0.5T MRI data. The table reports the mean and standard deviation of PSNR values for the reconstructed images generated by each method on the test set. The best performance is highlighted in red.

Datasets & Methods		Quantitative Evaluation PSNR
T2W Recon	0.5T ZF	19.92±0.79
	CycleGAN	19.65±1.22
	Score-MRI	20.17±1.03
	Meta-Score-MRI	21.22±1.25
T1W Recon	0.5T ZF	21.67±1.17
	CycleGAN	22.17±1.60
	Score-MRI	22.16±1.59
	Meta-Score-MRI	22.19±1.70

5.3. Generation of 0.5T-Like MR Image Form 1.5T MRI Data

Finally, we evaluate the generator G_{hl} trained with optimal transport (OT), which is designed to model the degradation mapping from 1.5T to 0.5T images and is essential for the effectiveness of the proposed method. Figure 8 shows real 0.5T images alongside pseudo 0.5T images generated from 1.5T data. Although both types of images contain considerable noise, the pseudo 0.5T images exhibit contrast characteristics that closely resemble those of the real 0.5T images. This indicates that G_{hl} is able to effectively capture the SNR and contrast degradation between 1.5T and 0.5T images.To further assess the fidelity of the generated pseudo 0.5T images, Table 5 reports the PSNR values between the generated images and the registered real 0.5T images in the test set. Despite the high noise levels present in both images, the average PSNR exceeds 20 dB, demonstrating that G_{hl} can accurately generate pseudo 0.5T images.

Table 5. Quantitative results for reconstructing 0.5T-like MR images from 1.5T MRI data. The table reports the mean and standard deviation of PSNR values for the reconstructed images generated by different methods on the test set.

Datasets	PSNR
T2W Recon	20.55±1.24
T1W Recon	20.84±1.66

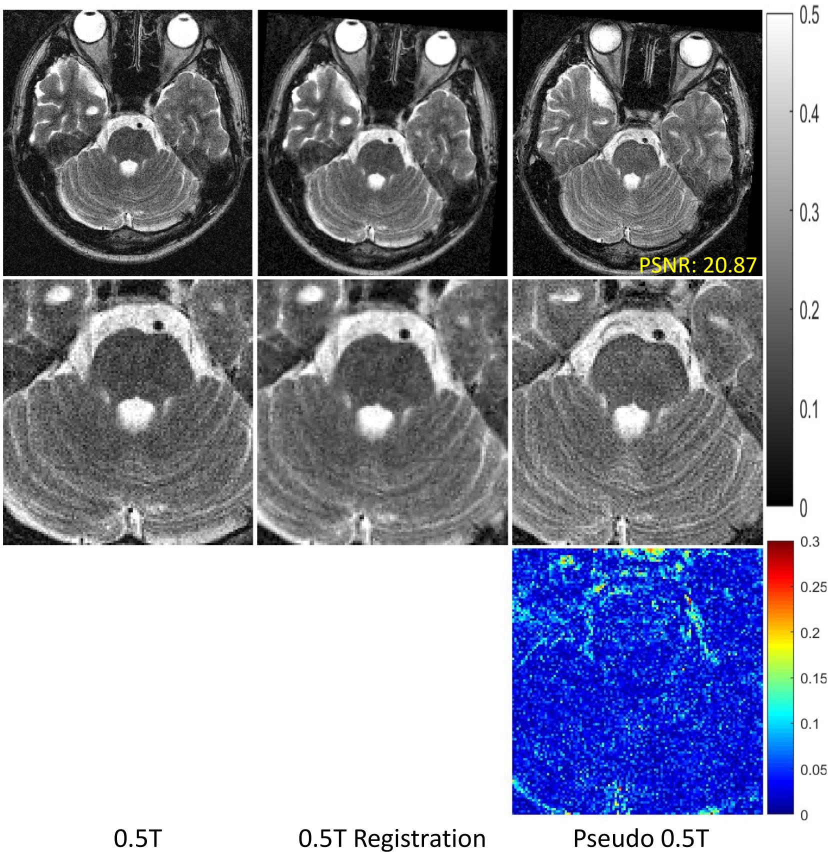

Figure 8. Demonstration of the capability of the OT-guided GAN (denoted as G_{hl}) to generate pseudo 0.5T-like images from 1.5T data. The first column shows the real T2-weighted images acquired from 0.5T MRI, while the second column presents the corresponding 0.5T images after registration with the 1.5T images. The third column displays the pseudo 0.5T-like T2-weighted images generated from 1.5T data. A zoomed-in view is provided in the second row, and the third row shows the error maps between the generated pseudo 0.5T-like images and the registered real 0.5T images. The grayscale for the reconstructed images and the color bar for the error maps are shown on the right side and in the middle of the figure, respectively.

6. Conclusions

This paper presents a model for reconstructing high-field-like MR images from low-field MRI, addressing the dual ill-posed problem arising from unpaired data. Both theoretical analysis and experimental validation indicate that conventional mapping-based learning approaches are insufficient for this task. To address this limitation, we propose a two-stage framework based on meta-learning that tackles the dual inverse problem from a regularization perspective. Experimental results demonstrate that the proposed method can reconstruct high-quality MR images from undersampled 0.5T k-space data, achieving SNR and contrast levels comparable to those of 1.5T MRI.

Low-field MRI provides a cost-effective and accessible imaging solution, particularly in resource-limited regions. However, its broader clinical adoption is constrained by lower spatial resolution and reduced SNR. The proposed method enables low-field MRI systems to produce images with quality approaching that of high-field MRI. Therefore, this approach has the potential to promote the wider deployment of MRI technology and generate meaningful societal benefits. The feasibility and effectiveness of the proposed method have been preliminarily validated using acquired unpaired low-field and high-field datasets. In future work, we will focus on collecting more extensive datasets to further evaluate the stability and robustness of the proposed framework.

Acknowledgments: The authors would like to express their gratitude to Shenzhen Anke High-tech Co., Ltd. for providing the 0.5T and 1.5T MRI data necessary for this paper, which were instrumental in validating the proposed methods. National Natural Science Foundation of China, Grant/Award Numbers: 61871373, 62271474, 81830056, 61771463, U1805261, 81729003, 81901736, 62106252, 81971611, and 62206273; Strategic Priority Research Program of Chinese Academy of Science, Grant/Award Numbers: XDB25000000, XDC07040000; Key Laboratory for Magnetic Resonance and Multimodality Imaging of Guangdong Province, Grant/Award Number: 2023B1212060052; Science and Technology Plan Program of Guangzhou, Grant/Award Number: 202007030002; Shenzhen Science and Technology Program,Grant/Award Numbers: KCXF20211020163408012, JCYJ20210324115810030.

References

1. Hori, M.; Hagiwara, A.; Goto, M.; Wada, A.; Aoki, S. Low-field magnetic resonance imaging: its history and renaissance. *Investigative Radiology* **2021**, *56*, 669.
2. Gecmen, D.B. Deep Learning Techniques for Low-Field MRI. Master's thesis, Delft University of Technology, 2020.
3. Tikhonov, A.N. On the stability of inverse problems. *Proceedings of the USSR Academy of Sciences* **1943**, *39*, 195–198.
4. Engl, H.W.; Hanke, M.; Neubauer, A. *Regularization of Inverse Problems*; Vol. 375, Springer Science & Business Media, 1996.
5. Cui, Z.X.; Zhu, Q.; Cheng, J.; Zhang, B.; Liang, D. Deep unfolding as iterative regularization for imaging inverse problems. *Inverse Problems* **2024**, *40*, 025011.
6. Bahrami, K.; Shi, F.; Zong, X.; Shin, H.W.; An, H.; Shen, D. Reconstruction of 7T-like images from 3T MRI. *IEEE Transactions on Medical Imaging* **2016**, *35*, 2085–2097.
7. Bahrami, K.; Shi, F.; Rekik, I.; Shen, D. Convolutional neural network for reconstruction of 7T-like images from 3T MRI using appearance and anatomical features. In Proceedings of the International Workshop on Deep Learning in Medical Image Analysis. Springer, 2016, pp. 39–47.
8. Wattjes, M.P.; Barkhof, F. High field MRI in the diagnosis of multiple sclerosis: high field–high yield? *Neuroradiology* **2009**, *51*, 279–292.
9. Van der Kolk, A.G.; Hendrikse, J.; Zwanenburg, J.J.; Visser, F.; Luijten, P.R. Clinical applications of 7 T MRI in the brain. *European journal of radiology* **2013**, *82*, 708–718.
10. Bahrami, K.; Shi, F.; Rekik, I.; Gao, Y.; Shen, D. 7T-guided super-resolution of 3T MRI. *Medical physics* **2017**, *44*, 1661–1677.
11. de Leeuw den Bouter, M.; Ippolito, G.; O'Reilly, T.; Remis, R.; van Gijzen, M.; Webb, A. Deep learning-based single image super-resolution for low-field MR brain images. *Scientific Reports* **2022**, *12*, 6362.
12. Bernstein, M.A. Field Strength Dependence in MRI: Advantages and Artifacts at 3T. In Proceedings of the Proceedings of the ISMRM, 2006.
13. Tamada, D.; Field, A.S.; Reeder, S.B. Simultaneous T1-weighted and T2-weighted 3D MRI using RF phase-modulated gradient echo imaging. *Magnetic Resonance in Medicine* **2022**, *87*, 1758–1770.
14. Adler, J.; Lunz, S. Banach Wasserstein GAN. In Proceedings of the Advances in Neural Information Processing Systems; Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; Garnett, R., Eds., 2018, Vol. 31.
15. Zhu, J.Y.; Park, T.; Isola, P.; Efros, A.A. Unpaired Image-To-Image Translation Using Cycle-Consistent Adversarial Networks. In Proceedings of the Proceedings of the IEEE International Conference on Computer Vision (ICCV), Oct 2017.
16. Sim, B.; Oh, G.; Kim, J.; Jung, C.; Ye, J.C. Optimal Transport Driven CycleGAN for Unsupervised Learning in Inverse Problems. *SIAM Journal on Imaging Sciences* **2020**, *13*, 2281–2306.
17. Oh, G.; Sim, B.; Chung, H.; Sunwoo, L.; Ye, J.C. Unpaired Deep Learning for Accelerated MRI Using Optimal Transport Driven CycleGAN. *IEEE Transactions on Computational Imaging* **2020**, *6*, 1285–1296.
18. Zou, L. *Meta-learning: Theory, Algorithms and Applications*; Elsevier Science & Technology, 2022.

19. Finn, C.; Abbeel, P.; Levine, S. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. In Proceedings of the Proceedings of the 34th International Conference on Machine Learning, 2017, pp. 1126–1135.
20. Zhu, M.; Li, S.; Li, D.; Gao, Q.; Bian, Z.; Huang, J.; Zeng, D.; Ma, J. Teacher-student Network for CT Image Reconstruction via Meta-learning Strategy. In Proceedings of the 2019 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), 2019, pp. 1–3.
21. Yang, L.; Zhang, Z.; Song, Y.; Hong, S.; Xu, R.; Zhao, Y.; Shao, Y.; Zhang, W.; Cui, B.; Yang, M.H. Diffusion models: A comprehensive survey of methods and applications. *arXiv preprint arXiv:2209.00796* **2022**.
22. Kazerouni, A.; Aghdam, E.K.; Heidari, M.; Azad, R.; Fayyaz, M.; Hacihaliloglu, I.; Merhof, D. Diffusion models for medical image analysis: A comprehensive survey. *arXiv preprint arXiv:2211.07804* **2022**.
23. Croitoru, F.A.; Hondru, V.; Ionescu, R.T.; Shah, M. Diffusion models in vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **2023**.
24. Song, Y.; Sohl-Dickstein, J.; Kingma, D.P.; Kumar, A.; Ermon, S.; Poole, B. Score-Based Generative Modeling through Stochastic Differential Equations. In Proceedings of the International Conference on Learning Representations, 2021.
25. Chung, H.; Ye, J.C. Score-based diffusion models for accelerated MRI. *Medical Image Analysis* **2022**, p. 102479.
26. Cao, C.; Cui, Z.X.; Wang, Y.; Liu, S.; Chen, T.; Zheng, H.; Liang, D.; Zhu, Y. High-Frequency Space Diffusion Model for Accelerated MRI. *IEEE Transactions on Medical Imaging* **2024**, pp. 1–1.
27. Cui, Z.X.; Cao, C.; Cheng, J.; Jia, S.; Zheng, H.; Liang, D.; Zhu, Y. SPIRiT-Diffusion: Self-Consistency Driven Diffusion Model for Accelerated MRI. *arXiv preprint arXiv:2304.05060* **2023**.
28. Ho, J.; Jain, A.; Abbeel, P. Denoising Diffusion Probabilistic Models. In Proceedings of the Advances in Neural Information Processing Systems, 2020, Vol. 33, pp. 6840–6851.
29. Vincent, P. A connection between score matching and denoising autoencoders. *Neural Computation* **2011**, *23*, 1661–1674.
30. Peyré, G.; Cuturi, M. Computational Optimal Transport: With Applications to Data Science. *Foundations and Trends® in Machine Learning* **2019**, *11*, 355–607.
31. Miyato, T.; Kataoka, T.; Koyama, M.; Yoshida, Y. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957* **2018**.
32. Song, Y.; Ermon, S. Improved Techniques for Training Score-Based Generative Models. In Proceedings of the Advances in Neural Information Processing Systems, 2020, Vol. 33, pp. 12438–12448.
33. Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. PyTorch: An Imperative Style, High-Performance Deep Learning Library, 2019, [arXiv:cs.LG/1912.01703].