

TRAINING PHYSICS-DRIVEN DEEP LEARNING RE-CONSTRUCTION WITHOUT RAW DATA ACCESS FOR EQUITABLE FAST MRI

Anonymous authors

Paper under double-blind review

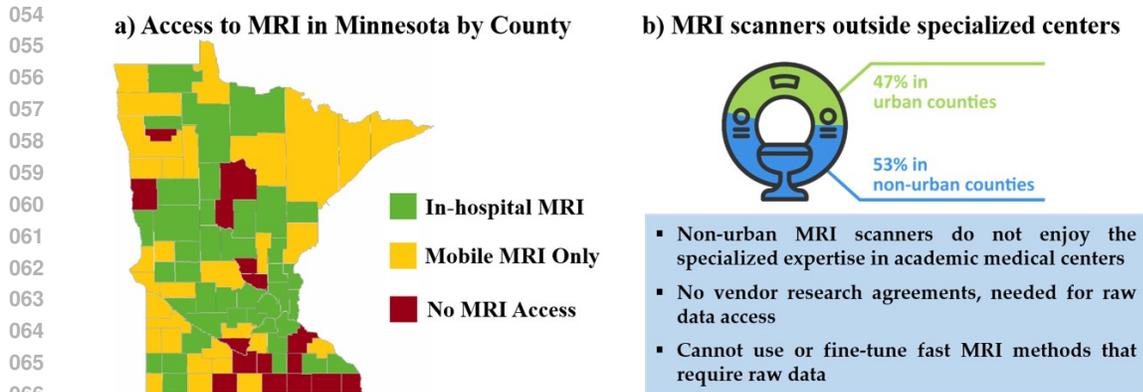
ABSTRACT

Physics-driven deep learning (PD-DL) approaches have become popular for improved reconstruction of fast magnetic resonance imaging (MRI) scans. Even though PD-DL offers higher acceleration rates compared to existing clinical fast MRI techniques, their use has been limited outside specialized MRI centers. One impediment for their deployment is the difficulties with generalization to pathologies or population groups that are not well-represented in training sets. This has been noted in several studies, and fine-tuning on target populations to improve reconstruction has been suggested. However, current approaches for PD-DL training require access to raw k-space measurements, which is typically only available at specialized MRI centers that have research agreements for such data access. This is especially an issue for rural and underserved areas, where commercial MRI scanners only provide access to a final reconstructed image. To tackle these challenges, we propose Compressibility-inspired Unsupervised Learning via Parallel Imaging Fidelity (CUPID) for high-quality PD-DL training using only routine clinical reconstructed images exported from an MRI scanner. CUPID evaluates the goodness of the output with a compressibility-based approach, while ensuring that the output stays consistent with the clinical parallel imaging reconstruction through well-designed perturbations. Our results show that CUPID achieves similar quality compared to well-established PD-DL training strategies that require raw k-space data access, while outperforming conventional compressed sensing (CS) and state-of-the-art generative methods. We also demonstrate its effectiveness in a zero-shot training setup for retrospectively and prospectively subsampled acquisitions, attesting to its minimal training burden. As an approach that radically deviates from existing strategies, CUPID presents an opportunity to provide equitable access to fast MRI for underserved populations in an attempt to reduce the inequalities associated with this expensive imaging modality.

1 INTRODUCTION

Magnetic resonance imaging (MRI) is a central tool in modern medicine, offering multiple soft tissue contrasts and high diagnostic sensitivity for numerous diseases. However, MRI is among the most expensive medical imaging modalities, in part due to its long scan times. Demand for MRI scans has shown an annual growth rate of 2.5%, while the number of MRI units per capita has increased by 1.8% in a similar time frame (Martella et al., 2023). This mismatch has further increased the wait times for MRI exams (Bartsch et al., 2024; Hofmann et al., 2023), particularly in rural areas and underserved communities (Burdorf, 2022), as depicted in Fig. 1. Thus, techniques for fast MRI scanning that can reduce overall scan times without compromising diagnostic quality (Akçakaya et al., 2022) are critical for improving the throughput of MRI.

Computational MRI approaches, including partial Fourier imaging (McGibney et al., 1993), parallel imaging (Pruessmann et al., 1999; Griswold et al., 2002), compressed sensing (Lustig et al., 2007), and more recently deep learning (DL) (Hammernik et al., 2018; Schlemper et al., 2018) have been developed for accelerating MRI. In most MRI centers, parallel imaging remains the most widely used approach for the reconstruction of routine clinical images. The acceleration rates afforded by these methods, however, are limited due to noise amplification and aliasing artifacts. DL-based



067
068
069
070
071
072

Figure 1: Many areas in the world have limited access to MRI or rely on services from local hospitals and mobile MRI: (a) A recent study from Minnesota, United States (Burdorf, 2022) show more than half of MRI services are in non-urban counties. (b) These non-urban MRI machines are unlikely to have vendor agreements that provide raw MRI data access, precluding access to and fine-tuning of AI-based fast MRI techniques.

073
074
075
076
077
078
079

methods, especially physics-driven DL (PD-DL) approaches, offer state-of-the-art improvements over parallel imaging (Knoll et al., 2020a). However, the translation of PD-DL to clinic has been hindered by generalizability and artifact issues related to details not well represented in training databases, in other words when faced with out-of-distribution samples at test time (Eldar et al., 2017; Knoll et al., 2020b; Muckley et al., 2021; Antun et al., 2020). This is a problem for many typical MRI centers, whose population characteristics do not necessarily align with specialized MRI centers in urban settings, where training databases for PD-DL are currently curated.

080
081
082
083
084
085

In such cases, fine-tuning of the PD-DL model on the target population may be beneficial (Knoll et al., 2019; Dar et al., 2020; Yaman et al., 2022b; Chandler et al., 2023). However, a major roadblock for this strategy is that all current training methods for PD-DL require access to raw MRI data. Such access requires research agreements with MR vendors, and is typically not available outside specialized/academic MRI centers. This is especially an issue for rural and underserved areas, where commercial MRI scanners only provide access to a final image, reconstructed via parallel imaging.

086
087
088
089
090
091
092

In this work, we tackle these challenges associated with typical PD-DL training, and propose Compressibility-inspired Unsupervised Learning via Parallel Imaging Fidelity (CUPID), which trains PD-DL reconstruction from routine clinical images, for instance in Digital Imaging and Communications in Medicine (DICOM) format. Succinctly, CUPID uses a compressibility-inspired term to evaluate the goodness of the output, while ensuring the output is consistent with parallel imaging via well-designed input perturbations. CUPID can be used both with database-training and in a subject-specific/zero-shot manner, attesting to its minimal fine-tuning burden.

093
094

Our key contributions include:

- 095
096
097
098
099
100
101
102
103
104
105
106
107
- We introduce CUPID a novel method that enables high-quality training of PD-DL reconstruction in unsupervised and zero-shot/subject-specific settings using only routine clinical reconstructed MRI images, eliminating the need for access to raw k-space measurements. To the best of our knowledge, our method is the first attempt to train PD-DL networks using only these images that are exported from the scanner.
 - CUPID trains on DICOM images acquired at the *target acceleration rate*, which often have noise and aliasing artifacts due to high sub-sampling, in an unsupervised manner. Note this is a deviation from other methods that use reference fully-sampled DICOM images to train a likelihood or score function, such as generative models.
 - CUPID uses a novel unsupervised loss formulation that enforces fidelity with using parallel imaging algorithms via carefully designed perturbations, in addition to evaluating the compressibility of the output image. This parallel imaging fidelity ensures the network does not converge to overly sparse solutions.

- We provide a comprehensive evaluation, encompassing acquisitions with both retrospective and prospective undersampling at target acceleration rates, to demonstrate that CUPID achieves results on par with leading supervised and self-supervised training strategies that depend on raw k-space data, while surpassing conventional compressed sensing (CS) techniques and state-of-the-art generative methods.
- CUPID enables training/fine-tuning of PD-DL reconstruction from routine clinically reconstructed DICOM images, and may be instrumental to provide equitable access to fast MRI methods, especially in underserved and rural areas.

2 BACKGROUND AND RELATED WORK

2.1 MRI FORWARD MODEL AND CONVENTIONAL METHODS FOR MRI RECONSTRUCTION

MRI raw data is acquired in the frequency domain of the image, referred to as k-space. For fast MRI, data is acquired in the sub-Nyquist regime by undersampling the acquisition in k-space. In this case, the forward acquisition model relating the image $\mathbf{x} \in \mathbb{C}^n$ to these raw MRI data (or k-space) measurements is given as:

$$\mathbf{y}_\Omega = \mathbf{E}_\Omega \mathbf{x} + \mathbf{n}, \quad (1)$$

where \mathbf{y}_Ω denotes the acquired k-space data corresponding to the undersampling pattern Ω with $|\Omega| = m < n$. \mathbf{E}_Ω denotes the multi-coil encoding operator that includes information from n_c receiver coils, each of which are sensitive to a different part of the image (Hamilton et al., 2017). When the acceleration rate $R = n/m$ is less than n_c , this system of equations is over-determined due to the redundancies among the receiver coils. Parallel imaging uses these redundancies to solve the maximum likelihood estimation problem under i.i.d. Gaussian noise (Pruessmann et al., 1999):

$$\mathbf{x}_{\text{PI}} = \arg \min_{\mathbf{x}} \|\mathbf{y}_\Omega - \mathbf{E}_\Omega \mathbf{x}\|_2^2 = (\mathbf{E}_\Omega^H \mathbf{E}_\Omega)^{-1} \mathbf{E}_\Omega^H \mathbf{y}_\Omega. \quad (2)$$

Numerically, this can be solved directly for certain undersampling patterns (Pruessmann et al., 1999) or more broadly iteratively using conjugate gradient (CG) (Pruessmann et al., 2001). Using the equivalence of multiplication in image domain and convolutions in k-space (Uecker et al., 2014), it can also be solved as an interpolation problem in k-space (Griswold et al., 2002). Parallel imaging remains the most clinically used acceleration method for MRI, with some MR systems using the image-based reconstruction, while others utilizing the equivalent k-space interpolation.

In modern computational MRI, additional regularization is often incorporated into the objective function (Hammernik et al., 2023):

$$\arg \min_{\mathbf{x}} \|\mathbf{y}_\Omega - \mathbf{E}_\Omega \mathbf{x}\|_2^2 + \mathcal{R}(\mathbf{x}), \quad (3)$$

where $\mathcal{R}(\cdot)$ denotes a regularizer. For instance, compressed sensing (CS) uses the idea that images should be compressible in an appropriate transform domain (Lustig et al., 2007), and uses $\mathcal{R}(\mathbf{x}) = \tau \|\mathbf{W}\mathbf{x}\|_1$, where τ is the regularization weight, \mathbf{W} is a linear sparsifying transform such as a discrete wavelet transform (DWT) and $\|\cdot\|_1$ is the ℓ_1 norm.

2.2 PD-DL RECONSTRUCTION VIA ALGORITHM UNROLLING

Among different PD-DL methods (Ahmad et al., 2020; Gilton et al., 2021; Knoll et al., 2020a), unrolled networks (Monga et al., 2021) remain the highest performer in reconstruction challenges, as reported a year ago (Hammernik et al., 2023; Muckley et al., 2021). These methods unroll iterative algorithms for solving the regularized least squares objective in (3) (Fessler, 2020), such as proximal gradient descent (Schlemper et al., 2018) or variable splitting with quadratic penalty (VS-QP) (Agarwal et al., 2019), over a fixed number of steps. VS-QP transforms (3) into 2 sub-problems:

$$\mathbf{z}^{(i)} = \arg \min_{\mathbf{z}} \|\mathbf{x}^{(i-1)} - \mathbf{z}\|_2^2 + \mathcal{R}(\mathbf{z}), \quad (4a)$$

$$\mathbf{x}^{(i)} = \arg \min_{\mathbf{x}} \|\mathbf{y}_\Omega - \mathbf{E}_\Omega \mathbf{x}\|_2^2 + \mu \|\mathbf{x} - \mathbf{z}^{(i)}\|_2^2, \quad (4b)$$

where (4a) is the proximal operator for the regularization, implicitly solved using neural networks, while (4b) accounts for data fidelity and has a closed form solution:

$$\mathbf{x}^{(i)} = (\mathbf{E}_\Omega^H \mathbf{E}_\Omega + \mu \mathbf{I})^{-1} (\mathbf{E}_\Omega^H \mathbf{y}_\Omega + \mu \mathbf{z}^{(i)}), \quad (5)$$

which can be solved by CG (Aggarwal et al., 2019). Unrolled networks are conventionally trained using supervised learning over a database, where the reference raw k-space measurements are first retrospectively undersampled to form \mathbf{y}_Ω . Subsequently, the network is trained to map to the original full reference k-space or the corresponding reference image (Hammernik et al., 2018; Aggarwal et al., 2019) by minimizing:

$$\min_{\theta} \mathbb{E} [\mathcal{L}(\mathbf{y}_{\text{ref}}, \mathbf{E}_{\text{full}}(f(\mathbf{y}_\Omega, \mathbf{E}_\Omega; \theta)))] \quad (6)$$

where θ are the network parameters, $f(\mathbf{y}_\Omega, \mathbf{E}_\Omega; \theta)$ denotes the network output for inputs \mathbf{y}_Ω and \mathbf{E}_Ω , \mathbf{E}_{full} is the fully-sampled encoding operator, \mathbf{y}_{ref} is the fully-sampled reference k-space data, and $\mathcal{L}(\cdot, \cdot)$ is a loss function.

2.3 SELF-SUPERVISED AND UNSUPERVISED METHODS

Obtaining fully-sampled reference data in MRI can be infeasible due to prolonged scan durations, organ movement in acquisitions such as real-time cardiac imaging or myocardial perfusion (Rajiah et al., 2023), or signal decay in acquisitions like diffusion MRI with EPI (Uğurbil et al., 2013). To enable training of PD-DL networks without fully sampled raw MRI data, a variety of unsupervised learning methodologies have emerged (Akçakaya et al., 2022), including self-supervised learning techniques (Yaman et al., 2020; Chen et al., 2021) and generative modeling approaches (Jalal et al., 2021; Chung & Ye, 2022; Chung et al., 2023).

Self-supervised methods use a masking approach to generate supervisory labels from the undersampled data (Yaman et al., 2020; Millard & Chiew, 2023; Hu et al., 2024). A pioneering method in this field, self-supervision via data undersampling (SSDU) (Yaman et al., 2020; 2022a), involves partitioning the acquired measurement Ω into two disjoint subsets ($\Omega = \Lambda \cup \Theta$) to train the network in a self-supervised manner:

$$\min_{\theta} \mathbb{E} [\mathcal{L}(\mathbf{y}_\Lambda, \mathbf{E}_\Lambda(f(\mathbf{y}_\Theta, \mathbf{E}_\Theta; \theta)))] \quad (7)$$

Even though these self-supervision based approaches demonstrate exceptional performance across various tasks, they lack the ability to train the model without access to undersampled raw data, as they cannot operate solely using images that are exported from the scanner.

Conversely, generative methods learn the prior distribution of the given dataset, which is then leveraged in conjunction with a log-likelihood data term during the testing phase. Although recent methods based on diffusion/score-based models have shown substantial promise, these methods require large amounts of high-quality images either reconstructed from raw data (Jalal et al., 2021; Luo et al., 2023) or as DICOMs (Chung & Ye, 2022), as well as computational resources to perform the training, both of which may not be feasible in the setups we are focused on.

3 UNSUPERVISED TRAINING FOR PD-DL WITHOUT RAW K-SPACE DATA

In this study, we introduce a novel framework to train PD-DL models, utilizing only routinely available clinical images exported directly from MRI scanners. Recently, inspired by the connections between PD-DL and compressibility-based processing (Gu et al., 2022), a compressibility-inspired loss was proposed to evaluate the goodness of unsupervised PD-DL training (Alçalar et al., 2024). However, this approach still requires access to raw k-space data to stabilize training, making it unsuitable for our goals. Here, we adapt the compressibility idea and augment it with a parallel imaging fidelity to successfully reconstruct clinical images in DICOM format without needing any raw k-space data.

Compressibility Aspect of the Loss Formulation. Compressibility/sparsity in the output of the PD-DL network can be enforced by utilizing a weighted ℓ_1 norm (Alçalar et al., 2024), which has been demonstrated to provide a closer approximation to the ℓ_0 norm compared to the standard ℓ_1 norm (Candes et al., 2008). Thus, this compressibility of the output image in CUPID is achieved by the loss term

$$\mathcal{L}_{\text{comp}}(\mathbf{x}_{\text{PI}}, \mathbf{m}) = \frac{1}{N} \cdot \sum_{n=1}^N \left(\frac{|(\mathbf{W}f(\mathbf{x}_{\text{PI}}, \mathbf{E}_\Omega))_n|}{|(\mathbf{W}\mathbf{x}^{(\mathbf{m})})_n| + \epsilon} \right), \quad (8)$$

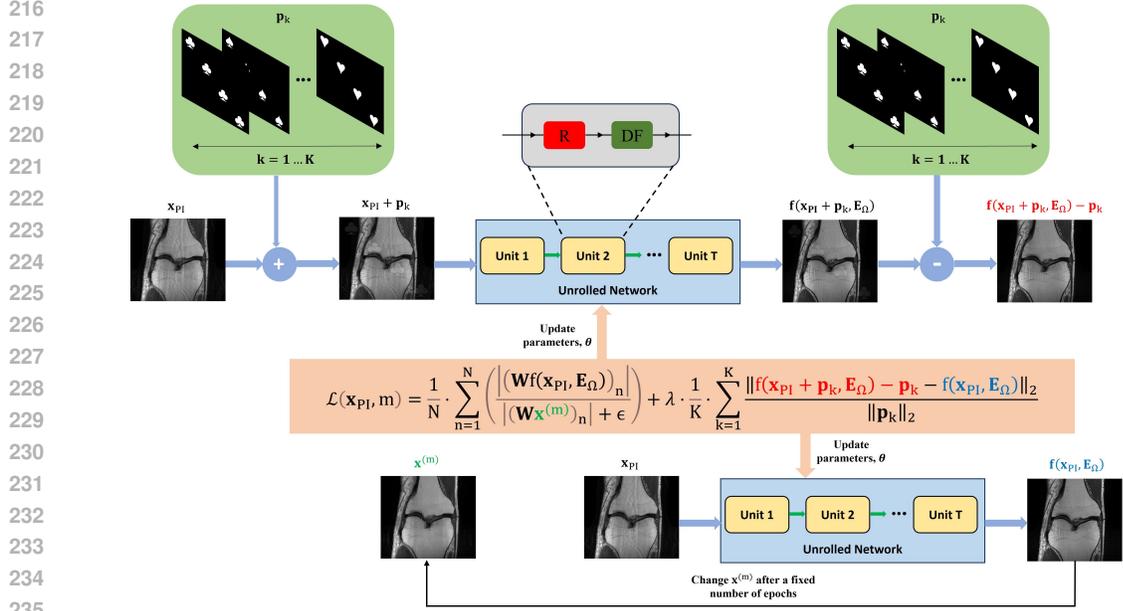


Figure 2: Our Compressibility-inspired Unsupervised Learning via Parallel Imaging Fidelity (CU-PID) method trains PD-DL models in an unsupervised and/or zero-shot manner without requiring any raw k -space data. The network is unrolled for T units, with each unit consisting of regularizer (R) and data fidelity (DF). The first term in the loss function is a reweighted ℓ_1 component designed to assess the compressibility of the network’s output, while the second term is a fidelity term that ensures the network does not produce a zero output by carefully perturbing the input data to check if it stays consistent with the parallel imaging reconstruction.

where x_{PI} denotes the DICOM input acquired using parallel imaging, \mathbf{W} represents the wavelet transform, N is the total number of wavelet coefficients and $x^{(m)}$ signifies the signal estimate following the training during the m^{th} reweighting step. Similar to Alçalar et al. (2024); Candes et al. (2008), we chose the initial weights from a CS reconstruction that has a large regularization and ϵ is added for numerical stability. Note, here we redefined $f(\cdot, \cdot)$ without the network parameters, θ , and used x_{PI} as the network input instead of y_{Ω} , to simplify notation.

Parallel Imaging Fidelity. Relying solely on (8) will result in inaccurate training as the neural network learns to produce an all-zeros image in an effort to drive the wavelet coefficients in the numerator to zero, which minimizes the loss function in (8). In Alçalar et al. (2024), fidelity with raw k -space data was used to avoid this training issue. In our setting without raw k -space access, we introduce a novel fidelity operator that stabilizes the training of the reconstruction algorithm, building on ideas from parallel imaging.

Specifically, we ensure that our network outputs are consistent with any clinical parallel imaging reconstruction through carefully crafted perturbations, $\{p_k\}$. These perturbations for R -fold acceleration are designed in such a way that R -fold aliasing do not create overlaps in the field-of-view, indicating that they could be resolved by parallel imaging reconstruction. The idea behind this design choice is to ensure that the network, when applied to the unperturbed x_{PI} , yields an accurate estimate of x , and when applied to $x_{PI} + p$, similarly recovers $x + p$, as the perturbation p must be resolvable within the framework of any parallel imaging approach. Both processes are visualized in Fig. 2. By doing so, the consistency term ensures a non-zero output when the sparsity is minimized. Thus, our second loss term that enforces parallel imaging fidelity is given as:

$$\mathcal{L}_{\text{pif}}(x_{PI}) = \mathbb{E}_{\mathbf{p}} \left[\frac{\|f(x_{PI} + \mathbf{p}, E_{\Omega}) - \mathbf{p} - f(x_{PI}, E_{\Omega})\|_2}{\|\mathbf{p}\|_2} \right]. \quad (9)$$

From an implementation perspective, the expectation over \mathbf{p} is calculated over K such perturbations $\{p_k\}$. The fold-over constraint for each $\{p_k\}$ is achieved by picking the perturbations as randomly rotated and positioned letters, numbers, card suits or other shapes that have different intensity

values. These choices also ensure that high-frequency information, such as edges, are accurately reconstructed by the regularization process. Our final loss function for CUPID is:

$$\mathcal{L}_{\text{CUPID}} = \mathcal{L}_{\text{comp}} + \lambda \cdot \mathcal{L}_{\text{pif}}, \quad (10)$$

where λ is a trade-off parameter between two terms.

Subject-Specific / Zero-Shot Application. In resource-limited or underserved settings, it may be more practical to fine-tune the method using only a few subjects, or even a single subject, to significantly reduce computational costs. As (8) does not solely focus on the subtraction between two entities, it lacks an inherent mechanism to drive the loss to zero through overfitting. Therefore, CUPID can be tailored to suit a scan-specific context (Akçakaya et al., 2019) without any modification to the loss given in (10).

4 EVALUATION

4.1 EXPERIMENTAL SETUP AND IMPLEMENTATION DETAILS

We conducted a thorough evaluation of our method, assessing its performance through both qualitative and quantitative analyses, and focused on uniform/equidistant patterns which produces coherent artifacts that are more difficult to remove compared to the incoherent artifacts from random undersampling (Knoll et al., 2019). We further note that CUPID demonstrates robust performance across a wide range of λ values, provided that λ is chosen within a reasonable range. An ablation study on the choice of λ is included in Sec. 4.5.

Retrospective Undersampling Setup. In our retrospective studies, we used fully-sampled multi-coil knee and brain MRI data from the fastMRI database (Knoll et al., 2020b). Knee dataset included fully-sampled coronal proton density-weighted (coronal PD) and PD with fat suppression (coronal PD-FS) data. For brain MRI, axial FLAIR (ax-FLAIR) dataset with matrix size of 320×320 is used. The knee and brain MRI datasets comprised data collected from 15 and 20 receiver coils, respectively. Both datasets were retrospectively undersampled using a uniform/equidistant pattern at $R = 4$. 24 lines of auto-calibration signal (ACS) from center of the raw k-space data were kept. DICOM images to train our proposed model were reconstructed using parallel imaging (CG-SENSE), solving $\mathbf{x}_{\text{PI}} = (\mathbf{E}_{\Omega}^H \mathbf{E}_{\Omega})^{-1} \mathbf{E}_{\Omega}^H \mathbf{y}_{\Omega}$. For each dataset, models were trained using 300 slices, and testing was performed using 380 slices for knee MRI and 100 slices for brain MRI, from distinct subjects.

Prospective Undersampling Setup. A multi-echo 3D GRE sequence on a 7T Siemens Magnetom MRI scanner was acquired. In this experiment, we replicate the practical pipeline for CUPID, where data is acquired at the desired high acceleration rate, and reconstructed to \mathbf{x}_{PI} with noise and aliasing artifacts, using parallel imaging. The corresponding DICOM images are exported and used for fine-tuning the PD-DL model with CUPID. To this end, the brain dataset, with matrix size = 288×288 and in-plane resolution $0.7 \times 0.7\text{mm}^2$, was acquired with prospective undersampling $R = 9$ (in k_y only), which is the desired target acceleration, much higher than the clinical protocol at $R = 3$. Low-resolution images were acquired in the same orientation for sensitivity estimation (Krueger et al., 2023). Training and reconstruction with CUPID was done in a zero-shot subject-specific manner.

More details about the implementation of the PD-DL models are provided in Appendix A.

4.2 COMPARISON METHODS

We compared our method against several database training methods that have access to raw k-space data, including supervised PD-DL (Hammernik et al., 2018; Aggarwal et al., 2019; Knoll et al., 2020a), self-supervision via data undersampling (SSDU) (Yaman et al., 2020), and equivariant imaging (EI) (Chen et al., 2021). All PD-DL methods utilized the same unrolled network and components (Appendix A) to ensure that only the training process differed for fair comparisons.

In addition, we compared our approach with methods that can operate without raw data access as long as \mathbf{E}_{Ω} is known at test time. These include compressed sensing (CS) (Lustig et al., 2007), and ScoreMRI (Chung & Ye, 2022). The latter trains a time-dependent score function using denoising score matching on a large dataset of reference fully-sampled images, and uses this score function

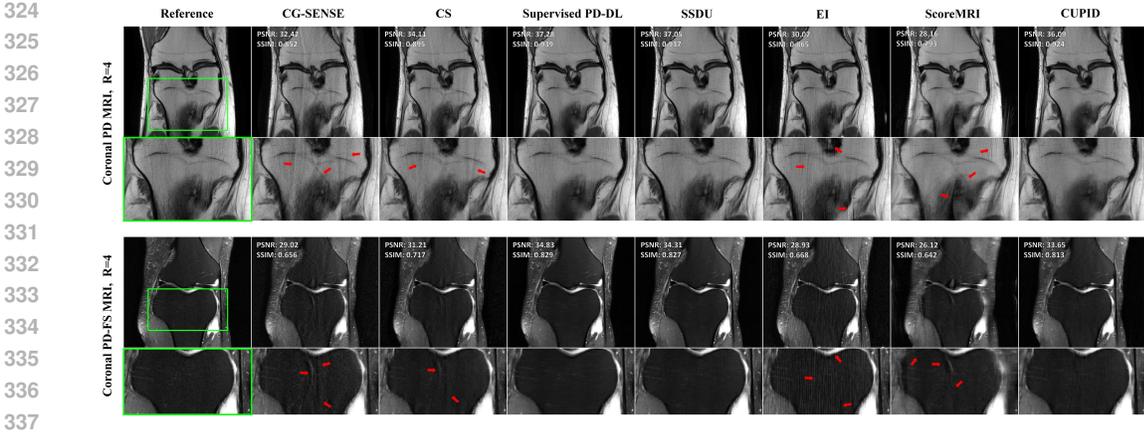


Figure 3: Representative coronal PD and PD-FS knee slices reconstructed using different methods for retrospective $R = 4$ uniform/equidistant undersampling. The baseline CG-SENSE, conventional CS, EI-trained PD-DL, and ScoreMRI suffers from residual artifacts highlighted by red arrows. PD-DL trained with CUPID loss function improves upon them while delivering reconstruction quality comparable to SSDU-trained PD-DL and supervised PD-DL.

during inference to sample from the conditional distribution given the measurements. Note both CS and ScoreMRI use $\mathbf{E}_\Omega^H \mathbf{y}_\Omega$ for data fidelity during inference. This can be accessed by multiplying \mathbf{x}_{PI} with $\mathbf{E}_\Omega^H \mathbf{E}_\Omega$. Note that \mathbf{E}_Ω includes information about the undersampling pattern Ω , which is completely known from the acquisition parameters, and coil sensitivities, which can be estimated from separate calibration scans in DICOM format (Krueger et al., 2023). A similar observation applies to the data fidelity in (5) for unrolled networks, thus they can be used for inference using only \mathbf{x}_{PI} and \mathbf{E}_Ω . We emphasize that what sets CUPID apart from other PD-DL strategies is that it is the only one that can train the unrolled network without using \mathbf{y}_Ω . Thus, without loss of generality, \mathbf{E}_Ω is known both at training and testing for all methods. Finally, we also used CG-SENSE, which was used to generate the original \mathbf{x}_{PI} as the clinical baseline comparison.

In the zero-shot setup, we compared our zero-shot results with zero-shot SSDU (ZS-SSDU) (Yaman et al., 2022b) as well as ScoreMRI, compressed sensing (CS), and our baseline method, CG-SENSE - all of which are compatible with zero-shot inference. All quantitative evaluations used structural similarity index (SSIM) and peak signal-to-noise ratio (PSNR).

4.3 EXPERIMENTS WITH RETROSPECTIVELY UNDERSAMPLED DATA

Database Results. Representative results in Fig. 3 show that baseline CG-SENSE, CS, EI and ScoreMRI reconstructions exhibit residual artifacts. In contrast, CUPID successfully eliminates

Table 1: Quantitative results for comparison methods on Coronal PD, Coronal PD-FS and Ax-FLAIR datasets using equispaced undersampling pattern at $R = 4$. Top 3 rows: Access to raw data (for training); Last 4 rows: No raw data access.

Method	Coronal PD		Coronal PD-FS		Ax-FLAIR	
	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑
Supervised PD-DL	40.95	0.964	35.89	0.859	36.69	0.926
SSDU	40.12	0.956	35.35	0.856	36.98	0.929
EI	33.29	0.919	29.86	0.704	35.48	0.908
ScoreMRI	32.84	0.812	28.18	0.684	28.17	0.774
CS	36.71	0.917	32.30	0.749	32.93	0.865
CG-SENSE	35.38	0.873	30.13	0.702	30.25	0.801
CUPID (ours)	39.28	0.951	34.71	0.840	36.02	0.920

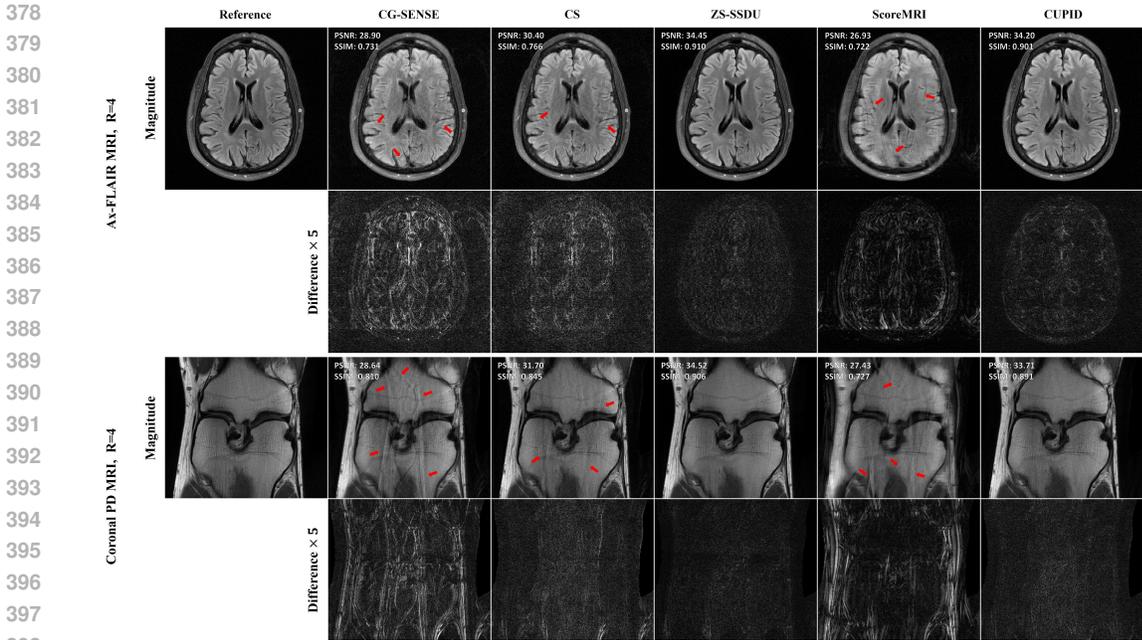


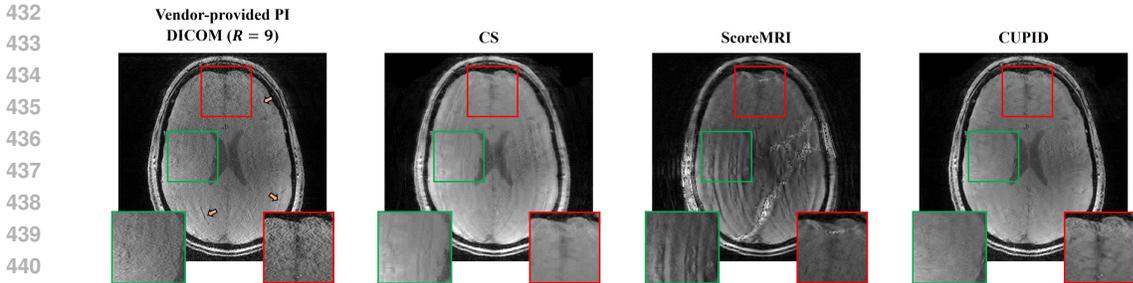
Figure 4: Representative subject-specific/zero-shot learning results for various algorithms on AX-FLAIR and coronal-PD dataset for retrospective $R = 4$ uniform undersampling. Baseline CG-SENSE, conventional CS and ScoreMRI suffer from residual artifacts (red arrows). PD-DL with CUPID loss successfully removes these artifacts, and functions in a similar manner to ZS-SSDU.

these artifacts from the CG-SENSE image using a well-trained PD-DL network, achieving state-of-the-art reconstruction quality comparable to supervised PD-DL and SSDU, despite only having access to x_{PI} for training, and not to raw k -space data unlike these other methods. We observe that parallel imaging reconstruction is not clinically usable at higher acceleration rates, but it is improved using a PD-DL reconstruction trained with CUPID. We further note that mild blurring was observed in some slices for database-training only. This is expected since we are no longer benefiting from redundancies from across multiple coils due to having no access to multi-coil raw k -space data, unlike the comparison methods. Quantitative results presented in Tab. 1 validate the visual observations, demonstrating that CUPID consistently outperforms CG-SENSE, CS, EI, and ScoreMRI across multiple datasets. Moreover, CUPID maintains performance comparable to that of supervised PD-DL and SSDU. Additional results on ax-FLAIR database are given in Appendix C.

Zero-Shot Learning Results. Fig. 4 shows results from zero-shot reconstructions. Both CG-SENSE and CS suffer from noise amplification and persistent residual artifacts, with CG-SENSE displaying a more pronounced degradation in quality. On the other hand, ScoreMRI exhibits blurring while still displaying residual artifacts. We note that uniform undersampling is used in these datasets, which is consistent with clinical parallel imaging acquisitions, but which have not been previously reported with diffusion models in existing works. Once again, CUPID demonstrates superior artifact and noise reduction over these methods and closely matches the quality of ZS-SSDU, despite not having access to raw data and an explicit self-validation mechanism to prevent overfitting as in the latter.

4.4 PRACTICAL SETTING: PROSPECTIVELY UNDERSAMPLED STUDY

As discussed in Sec. 4.1, brain data is acquired at the target acceleration rate, reconstructed via parallel imaging and exported in DICOM format to perform zero-shot fine-tuning. Fig. 5 shows reconstruction results for the vendor parallel imaging reconstruction, as well as CS, ScoreMRI and CUPID. CS reduces the noise in the parallel imaging reconstruction, but leads to blurring due to over-regularization. In contrast, ScoreMRI struggles to reconstruct accurately at this high acceleration rate, suggesting generalizability issues for the pre-trained score function to high-resolution imaging at a different field strength, not represented in the training database, and potential difficulties

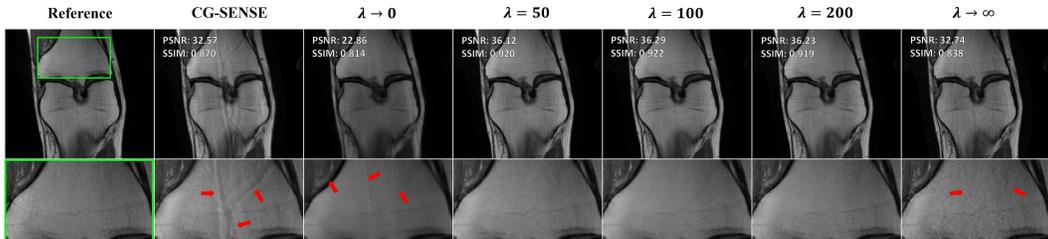


442 Figure 5: Prospective acceleration results for various methods that can operate on parallel imaging reconstructed DICOM images exported from the scanner. CUPID provides substantial noise and
443 artifact reduction on the DICOM image, outperforming other methods.
444

445
446
447 with uniform undersampling. Furthermore, the public implementation of ScoreMRI uses hard constraints for data fidelity, which leads to more pronounced artifacts due to potential phase mismatch with the coils generated from separate calibration data. Our proposed CUPID method effectively mitigates both artifacts and noise in the DICOM image (shown in zoomed insets) without requiring any raw k-space data, attesting to the effectiveness of CUPID in real-world scenarios. Note minor residual artifacts remain since the target acceleration $R = 9$ in 1-dimension is very high. We note that ZS-SSDU cannot be applied here due to the unavailability of raw data. We further note that the vendor-provided DICOM was generated using k-space interpolation (Griswold et al., 2002) instead of the image domain formulation in (2). Due to their equivalence, this did not cause any issues for CUPID, as expected.
448
449
450
451
452
453
454
455
456

457
458 **4.5 ABLATION STUDIES**
459

460 We carried out two ablation studies to explore key factors influencing the performance of our algorithm. The first study explored the effect of λ parameter to the final reconstruction, by training 5 distinct PD-DL networks using $\lambda \in \{0, 50, 100, 200, \infty\}$. We note that using $\lambda = 0$ corresponds to using only the compressibility term ($\mathcal{L}_{\text{comp}}$ in (8)), whereas using $\lambda \rightarrow \infty$ translates to using solely the parallel imaging fidelity term (\mathcal{L}_{pir} in (9)). Fig. 6 shows the corresponding reconstruction results for each case. As outlined in Sec. 3, only using $\mathcal{L}_{\text{comp}}$ leads to overly-smooth reconstructions due to network forcing the wavelet coefficients towards zero without maintaining consistency with the data.
461
462
463
464
465
466
467 On the other hand, solely using \mathcal{L}_{pir} results in DIP-like reconstructions (Ulyanov et al., 2018), where the network overfits the data without any regularization, resulting in noise amplification. CUPID with $\lambda \in \{50, 100, 200\}$ integrates both loss terms to attain high-fidelity reconstructions. Thus, we conclude that CUPID demonstrates robust performance across a wide range of λ values. Our second ablation study focuses on the effect of the number of perturbation patterns used in the training, and is provided in Appendix B.1.
468
469
470
471
472
473
474



483 Figure 6: Using only the compressibility term ($\lambda = 0$) in the loss leads to overly-smoothed images, whereas using only the parallel imaging fidelity term ($\lambda \rightarrow \infty$) causes noise amplification (red arrows). Using both terms with a mid-range λ value as a trade-off provides high-quality reconstructions that are clear from noise and artifacts.
484
485

4.6 DISCUSSION AND FUTURE WORK

Filtering on Routine Clinical Images. MR scans may include filtering operations applied by some vendors that affect the assumption $\hat{\mathbf{x}}_{\text{PI}} = (\mathbf{E}_{\Omega}^H \mathbf{E}_{\Omega})^{-1} \mathbf{E}_{\Omega}^H \mathbf{y}_{\Omega}$. This was discussed extensively in Shimron et al. (2022), in the context of using retrospective undersampling of DICOM images to train DL reconstruction, especially highlighting the use of zero-padding, which improves the display resolution compared to the acquisition resolution. It was shown that training of models from retrospective undersampling of databases of DICOM images for PD-DL training using zero-padding may lead to biases and inaccuracies. Conversely, our approach is physics-driven in nature, and the sampling pattern Ω naturally accounts for the zero-padding operation. However, our method is not immune to other types of filtering/processing, such as implicit intensity correction (Han et al., 2001) or deidentification methods (Van Essen et al., 2013), in which case the filtered \mathbf{x}_{PI} would need to be treated as the parallel imaging solution corresponding to a filtered version of \mathbf{y} .

Resources for Fine-Tuning of PD-DL Reconstruction. While our method is aimed to improve equitable access to fast MRI in low-resource settings, we do acknowledge that such low-resource MRI centers may lack the necessary hardware to fine-tune PD-DL models. We note that, similar to what has been shown in Yaman et al. (2022b), transfer learning along with zero-shot fine-tuning may be beneficial in decreasing training time and resources. Furthermore, since only DICOM images are needed for CUPID, standard anonymization techniques can be used (Van Essen et al., 2013), and data can be transferred to potential offsite locations with more computational resources. Since no raw data storage or transfer is necessary, the burden for transmission of DICOMs is low, and privacy can be preserved through anonymization, though specifics would need to be implemented with proper protocols.

Magnitude-only DICOMs. In some cases, MR vendors only give access to the magnitude of \mathbf{x}_{PI} due to limitations in the reconstruction pipelines, for instance when using partial Fourier imaging in some vendors. Additionally, if a reference calibration scan is not prospectively acquired at exam time, coil sensitivities, and thus \mathbf{E}_{Ω} would be unknown as well, which is the case for most existing DICOM databases. In such cases, the phase of \mathbf{x}_{PI} and coil sensitivities need to be estimated jointly. There has been work on the latter in the standard PD-DL setup with access to raw k-space data (Hu et al., 2024; Arvinte et al., 2021). However, this is a modification on the PD-DL objective function, and not the learning procedure considered here. Thus, it is not the focus of our study and it will be investigated in future studies.

Reinforcement Learning for Optimal Perturbations. Finally, a promising direction for future work may involve using reinforcement learning to simultaneously learn the optimal perturbations. This approach may enable the model to adapt and optimize perturbation strategies so that the sample mean estimate in (9) can be better approximated with a few perturbation examples.

5 CONCLUSION

In this study, we presented a novel training strategy, Compressibility-inspired Unsupervised Learning via Parallel Imaging Fidelity (CUPID), for PD-DL MRI reconstruction without access to raw k-space data. This approach leverages the compressibility of output images along with strategically designed perturbations that remain intact post-parallel imaging, thereby enhancing image quality in clinically accessible images in a physics-driven manner without the need for any raw k-space data. To the best of our knowledge, this is the first attempt that does not rely on raw data and uses these clinical images to train PD-DL networks, which is known for their high-fidelity reconstructions. CUPID also alleviates the training burden of generative methods, which requires a large number of data during training to capture the prior well. Quantitative and qualitative assessments of our method, conducted on both retrospectively and prospectively accelerated acquisitions, show its effectiveness in delivering high-quality performance across a diverse range of MRI scans and learning settings. Our source code will be released upon the acceptance of the paper.

REFERENCES

H. K. Aggarwal, M. P. Mani, and M. Jacob. MoDL: Model-Based Deep Learning Architecture for Inverse Problems. *IEEE Trans Med Imaging*, 38(2):394–405, Feb 2019.

- 540 R. Ahmad, C. A. Bouman, G. T. Buzzard, S. Chan, S. Liu, E. T. Reehorst, and P. Schniter. Plug-and-
541 play methods for magnetic resonance imaging: Using denoisers for image recovery. *IEEE signal*
542 *processing magazine*, 37(1):105–116, 2020.
- 543
- 544 M. Akçakaya, S. Moeller, S. Weingärtner, and K. Uğurbil. Scan-specific robust artificial-neural-
545 networks for k-space interpolation (RAKI) reconstruction: Database-free deep learning for fast
546 imaging. *Magn Reson Med*, 81:439–453, 2019.
- 547 M. Akçakaya, B. Yaman, H. Chung, and J. C. Ye. Unsupervised Deep Learning Methods for Biolog-
548 ical Image Reconstruction and Enhancement: An overview from a signal processing perspective.
549 *IEEE Sig Proc Mag*, 39(2):28–44, Mar 2022.
- 550
- 551 M. Akçakaya, M. Doneva, and C. P. (eds.). *Magnetic Resonance Image Reconstruction*, volume 7
552 of *Advances in Magnetic Resonance Technology and Applications*. Academic Press, 2022.
- 553
- 554 Y. Alçalar, M. Gülle, and M. Akçakaya. A convex compressibility-inspired unsupervised loss func-
555 tion for physics-driven deep learning reconstruction. In *2024 IEEE International Symposium on*
556 *Biomedical Imaging (ISBI)*, pp. 1–5. IEEE, 2024.
- 557
- 558 V. Antun, F. Renna, C. Poon, B. Adcock, and A. C. Hansen. On instabilities of deep learning in
559 image reconstruction and the potential costs of AI. *Proceedings of the National Academy of*
560 *Sciences*, 117(48):30088–30095, 2020.
- 561
- 562 M. Arvinte, S. Vishwanath, A. H. Tewfik, and J. I. Tamir. Deep J-Sense: Accelerated MRI re-
563 construction via unrolled alternating optimization. In *International conference on medical image*
564 *computing and computer-assisted intervention*, pp. 350–360. Springer, 2021.
- 565
- 566 E. Bartsch, S. Shin, K. Sheehan, M. Fralick, A. Verma, F. Razak, and L. Lapointe-Shaw. Advanced
567 imaging use and delays among inpatients with psychiatric comorbidity. *Brain Behav*, 14(2):
568 e3425, Feb 2024.
- 569
- 570 B. T. Burdorf. Comparing magnetic resonance imaging and computed tomography machine ac-
571 cessibility among urban and rural county hospitals. *Journal of Public Health Research*, 11(1):
572 jphr–2021, 2022.
- 573
- 574 E. J. Candes, M. B. Wakin, and S. P. Boyd. Enhancing sparsity by reweighted ℓ_1 minimization. *J*
575 *Fourier Anal Appl*, 14:877–905, 2008.
- 576
- 577 E. P. Chandler, S. Shoushtari, J. Liu, M. S. Asif, and U. S. Kamilov. Overcoming distribution shifts
578 in plug-and-play methods with test-time training. In *2023 IEEE 9th International Workshop on*
579 *Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, pp. 186–190. IEEE,
580 2023.
- 581
- 582 D. Chen, J. Tachella, and M. E. Davies. Equivariant imaging: Learning beyond the range space. In
583 *Proc IEEE/CVF ICCV*, pp. 4379–4388, 2021.
- 584
- 585 H. Chung and J. C. Ye. Score-based diffusion models for accelerated MRI. *Med Image Anal*, 80:
586 102479, Aug 2022.
- 587
- 588 H. Chung, S. Lee, and J. C. Ye. Decomposed diffusion sampler for accelerating large-scale inverse
589 problems. *arXiv preprint arXiv:2303.05754*, 2023.
- 590
- 591 F. Cotter. *Uses of complex wavelets in deep convolutional neural networks*. PhD thesis, University
592 of Cambridge, 2020.
- 593
- 594 S. U. H. Dar, M. Özbey, A. B. Çatlı, and T. Çukur. A transfer-learning approach for accelerated
595 MRI using deep neural networks. *Magnetic Resonance in Medicine*, 84(2):663–685, 2020.
- 596
- 597 Y. C. Eldar, A. O. Hero III, L. Deng, J. Fessler, J. Kovacevic, H. V. Poor, and S. Young. Challenges
598 and open problems in signal processing: Panel discussion summary from ICASSP 2017 [panel
599 and forum]. *IEEE Signal Processing Magazine*, 34(6):8–23, 2017. doi: 10.1109/MSP.2017.
2743842.

- 594 J. A. Fessler. Optimization methods for magnetic resonance image reconstruction. *IEEE Sig Proc*
595 *Mag*, 37(1):33–40, 2020.
- 596
- 597 D. Gilton, G. Ongie, and R. Willett. Deep equilibrium architectures for inverse problems in imaging.
598 *IEEE Transactions on Computational Imaging*, 7:1123–1133, 2021.
- 599
- 600 M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and
601 A. Haase. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magnetic Reso-*
602 *nance in Medicine*, 47(6):1202–1210, 2002.
- 603 H. Gu, B. Yaman, S. Moeller, J. Ellermann, K. Ugurbil, and M. Akçakaya. Revisiting ℓ_1 -wavelet
604 compressed-sensing MRI in the era of deep learning. *Proc Natl Acad Sci*, 119(33):e2201062119,
605 Aug 2022.
- 606
- 607 J. Hamilton, D. Franson, and N. Seiberlich. Recent advances in parallel imaging for MRI. *Progress*
608 *in nuclear magnetic resonance spectroscopy*, 101:71–95, 2017.
- 609
- 610 K. Hammernik, T. Klatzer, E. Kobler, M. P. Recht, D. K. Sodickson, T. Pock, and F. Knoll. Learning
611 a variational network for reconstruction of accelerated MRI data. *Magn Reson Med*, 79(6):3055–
612 3071, Jun 2018.
- 613
- 614 K. Hammernik, T. Küstner, B. Yaman, Z. Huang, D. Rueckert, F. Knoll, and M. Akçakaya. Physics-
615 driven deep learning for computational magnetic resonance imaging. *IEEE Sig Proc Mag*, 40:
616 98–114, 2023.
- 617
- 618 C. Han, T. S. Hatsukami, and C. Yuan. A multi-scale method for automatic correction of intensity
619 non-uniformity in MR images. *Journal of Magnetic Resonance Imaging.*, 13(3):428–436, 2001.
- 620
- 621 B. Hofmann, O. Brandsaeter, and E. Kjelle. Variations in wait times for imaging services: a register-
622 based study of self-reported wait times for specific examinations in Norway. *BMC Health Serv*
623 *Res*, 23(1):1287, Nov 2023.
- 624
- 625 Y. Hu, W. Gan, C. Ying, T. Wang, C. Eldeniz, J. Liu, Y. Chen, H. An, and U. S. Kamilov. SPICER:
626 Self-supervised learning for MRI with automatic coil sensitivity estimation and reconstruction.
627 *Magnetic Resonance in Medicine*, 2024.
- 628
- 629 A. Jalal, M. Arvinte, G. Daras, E. Price, A. G. Dimakis, and J. Tamir. Robust compressed sensing
630 MRI with deep generative priors. *Adv Neural Inf Process*, 34:14938–14954, 2021.
- 631
- 632 F. Knoll, K. Hammernik, C. Zhang, S. Moeller, T. Pock, D. K. Sodickson, and M. akaya. Deep-
633 Learning Methods for Parallel Magnetic Resonance Imaging Reconstruction: A Survey of the
634 Current Approaches, Trends, and Issues. *IEEE Sig Proc Mag*, 37(1):128–140, Jan 2020a.
- 635
- 636 F. Knoll, K. Hammernik, E. Kobler, T. Pock, M. P. Recht, and D. K. Sodickson. Assessment of the
637 generalization of learned image reconstruction and the potential for transfer learning. *Magnetic*
638 *Resonance in Medicine*, 81(1):116–128, 2019.
- 639
- 640 F. Knoll, J. Zbontar, A. Sriram, M. J. Muckley, M. Bruno, A. Defazio, M. Parente, K. J. Geras,
641 J. Katsnelson, H. Chandarana, et al. fastMRI: A publicly available raw k-space and DICOM
642 dataset of knee images for accelerated mr image reconstruction using machine learning. *Radiol-*
643 *ogy: Artificial Intelligence*, 2(1):e190007, 2020b.
- 644
- 645 F. Krueger, C. S. Aigner, K. Hammernik, S. Dietrich, M. Lutz, J. Schulz-Menger, T. Schaeffter, and
646 S. Schmitter. Rapid estimation of 2D relative B1+-maps from localizers in the human heart at 7T
647 using deep learning. *Magnetic Resonance in Medicine*, 89(3):1002–1015, 2023.
- 648
- 649 G. Luo, M. Blumenthal, M. Heide, and M. Uecker. Bayesian MRI reconstruction with joint un-
650 certainty estimation using diffusion models. *Magnetic Resonance in Medicine*, 90(1):295–311,
651 2023.
- 652
- 653 M. Lustig, D. Donoho, and J. M. Pauly. Sparse MRI: The application of compressed sensing for
654 rapid MR imaging. *Magn Reson Med*, 58(6):1182–1195, Dec 2007.

- 648 M. Martella, J. Lenzi, and M. M. Gianino. Diagnostic technology: Trends of use and availability
649 in a 10-year period (2011-2020) among sixteen OECD countries. *Healthcare (Basel)*, 11(14), Jul
650 2023.
- 651 G. McGibney, M. Smith, S. Nichols, and A. Crawley. Quantitative evaluation of several partial
652 Fourier reconstruction algorithms used in MRI. *Magnetic Resonance in Medicine*, 30(1):51–59,
653 1993.
- 654 C. Millard and M. Chiew. A theoretical framework for self-supervised MR image reconstruction
655 using sub-sampling via variable density Noisier2Noise. *IEEE Transactions on Computational*
656 *Imaging*, 2023.
- 657 V. Monga, Y. Li, and Y. C. Eldar. Algorithm unrolling: Interpretable, efficient deep learning for
658 signal and image processing. *IEEE Signal Processing Magazine*, 38(2):18–44, 2021.
- 659 M. J. Muckley, B. Riemenschneider, A. Radmanesh, S. Kim, G. Jeong, J. Ko, Y. Jun, H. Shin,
660 D. Hwang, M. Mostapha, et al. Results of the 2020 fastMRI challenge for machine learning MR
661 image reconstruction. *IEEE Transactions on Medical Imaging*, 40(9):2306–2317, 2021.
- 662 K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger. SENSE: sensitivity encoding for
663 fast MRI. *Magnetic Resonance in Medicine*, 42(5):952–962, 1999.
- 664 K. P. Pruessmann, M. Weiger, P. Börnert, and P. Boesiger. Advances in sensitivity encoding with
665 arbitrary k-space trajectories. *Magnetic Resonance in Medicine*, 46(4):638–651, 2001.
- 666 P. S. Rajiah, C. J. François, and T. Leiner. Cardiac MRI: State of the art. *Radiology*, 307(3):e223008,
667 May 2023.
- 668 J. Schlemper, J. Caballero, J. V. Hajnal, A. N. Price, and D. Rueckert. A deep cascade of convo-
669 lutional neural networks for dynamic mr image reconstruction. *IEEE transactions on Medical*
670 *Imaging*, 37(2):491–503, 2018.
- 671 I. W. Selesnick, R. G. Baraniuk, and N. C. Kingsbury. The dual-tree complex wavelet transform.
672 *IEEE Signal Processing Magazine*, 22(6):123–151, 2005.
- 673 E. Shimron, J. I. Tamir, K. Wang, and M. Lustig. Implicit data crimes: Machine learning bias
674 arising from misuse of public data. *Proceedings of the National Academy of Sciences*, 119(13):
675 e2117203119, 2022.
- 676 R. Timofte, E. Agustsson, L. Van Gool, M.-H. Yang, and L. Zhang. Ntire 2017 challenge on single
677 image super-resolution: Methods and results. In *Proceedings of the IEEE conference on computer*
678 *vision and pattern recognition workshops*, pp. 114–125, 2017.
- 679 M. Uecker, P. Lai, M. J. Murphy, P. Virtue, M. Elad, J. M. Pauly, S. S. Vasanawala, and
680 M. Lustig. ESPIRiT—an eigenvalue approach to autocalibrating parallel MRI: where SENSE
681 meets GRAPPA. *Mag Reson Med*, 71:990–1001, 2014.
- 682 K. Uğurbil, J. Xu, E. J. Auerbach, S. Moeller, A. T. Vu, J. M. Duarte-Carvajalino, C. Lenglet, X. Wu,
683 S. Schmitter, P. F. Van de Moortele, et al. Pushing spatial and temporal resolution for functional
684 and diffusion MRI in the Human Connectome Project. *Neuroimage*, 80:80–104, 2013.
- 685 D. Ulyanov, A. Vedaldi, and V. Lempitsky. Deep image prior. In *Proceedings of the IEEE conference*
686 *on computer vision and pattern recognition*, pp. 9446–9454, 2018.
- 687 D. C. Van Essen, S. M. Smith, D. M. Barch, T. E. Behrens, E. Yacoub, K. Ugurbil, W.-M. H.
688 Consortium, et al. The WU-Minn human connectome project: an overview. *Neuroimage*, 80:
689 62–79, 2013.
- 690 B. Yaman, S. A. H. Hosseini, S. Moeller, J. Ellermann, K. Uğurbil, and M. Akçakaya. Self-
691 supervised learning of physics-guided reconstruction neural networks without fully sampled refer-
692 ence data. *Magn Reson Med*, 84(6):3172–3191, Dec 2020.
- 693 B. Yaman, H. Gu, S. A. H. Hosseini, O. B. Demirel, S. Moeller, J. Ellermann, K. Uğurbil, and
694 M. Akçakaya. Multi-mask self-supervised learning for physics-guided neural networks in highly
695 accelerated magnetic resonance imaging. *NMR Biomed*, 35(12):e4798, Dec 2022a.

702 B. Yaman, S. A. H. Hosseini, and M. Akcakaya. Zero-shot self-supervised learning for MRI re-
703 construction. In *International Conference on Learning Representations, 2022b*. URL <https://openreview.net/forum?id=085y6YPaYjP>.
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755

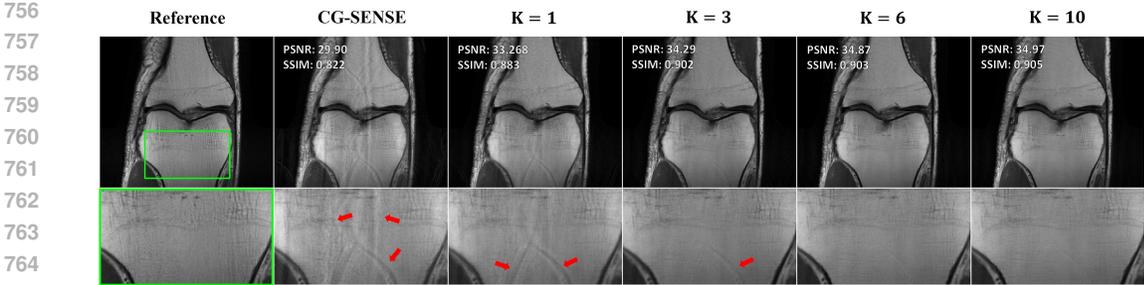


Figure 7: Representative zero-shot fine-tuning result for CUPID with different number of perturbations, $K \in \{1, 3, 6, 10\}$, on coronal PD knee MRI using $R = 4$ uniform undersampling. For $K < 6$, expectation in the second parallel imaging fidelity loss term can not converge to true mean, leading to artifacts in the final reconstruction. However, as K increases, reconstruction quality also enhances. Note the gains from further increasing K becomes negligible for $K > 6$.

A IMPLEMENTATION DETAILS FOR EACH METHOD

Compressed Sensing. We solved the regularized ℓ_1 minimization problem given below:

$$\arg \min_{\mathbf{x}} \|\mathbf{y}_\Omega - \mathbf{E}_\Omega \mathbf{x}\|_2 + \tau \|\mathbf{W} \mathbf{x}\|_1, \tag{11}$$

using VS-QP (Fessler, 2020). Similar to the unrolled network, data fidelity was solved using CG, and soft thresholding was implemented on the DTCWT coefficients.

ScoreMRI. For ScoreMRI implementation, we followed the original code and pre-trained network provided by Chung & Ye (2022) in their corresponding public repository.

PD-DL Based Approaches. For each method, the unrolled network comprised 10 unrolls, while the regularizer was implemented as a CNN-based ResNet architecture (Timofte et al., 2017) that had 10 residual blocks. Data fidelity was achieved using a conjugate-gradient (CG) method (Aggarwal et al., 2019) with 10 iterations. The unrolled network was trained in an end-to-end fashion for 100 epochs. For supervised PD-DL (Hammernik et al., 2018; Aggarwal et al., 2019), the normalized ℓ_1 - ℓ_2 loss function was used between the reconstructed and ground truth raw k-space data (Knoll et al., 2020a). For SSDU, $\rho = |\Delta|/|\Omega| = 0.4$ was used as proposed in Yaman et al. (2020). For EI (Chen et al., 2021), we modified the loss function in PD-DL networks to:

$$\min_{\theta} \mathbb{E} [\mathcal{L}(\mathbf{y}_\Omega, f(\mathbf{y}_\Omega, \mathbf{E}_\Omega; \theta))] + \beta \sum_{g \in G} \mathcal{L}(\mathcal{T}_g f(\mathbf{y}_\Omega, \mathbf{E}_\Omega; \theta), f(\mathbf{E}_\Omega \mathcal{T}_g f(\mathbf{y}_\Omega, \mathbf{E}_\Omega; \theta), \mathbf{E}_\Omega; \theta)) \tag{12}$$

in which the first term enforces consistency while the second term imposes equivariance relative to a group of transformations, $\{\mathcal{T}_g\}_{g \in G}$. Here, $|G|$ defined as the cardinality of $\{\mathcal{T}_g\}_{g \in G}$ and β is the equivariance weight. We followed the authors' publicly available CT reconstruction code

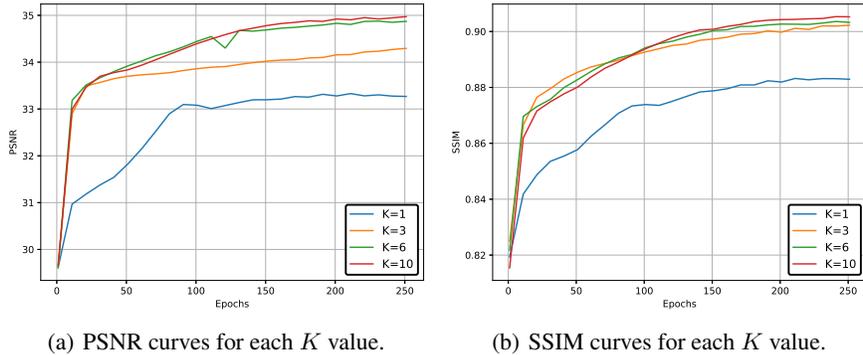


Figure 8: PSNR and SSIM curves confirm the visual observations with respect to the number of perturbations, K . Lower K values tend to perform worse and increasing K becomes redundant after a certain point.

810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863

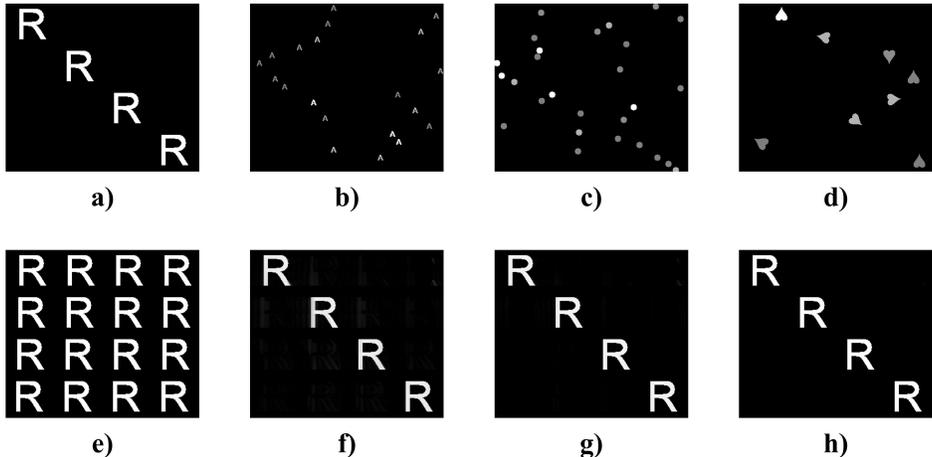


Figure 9: Added perturbations may consist of: (a) precisely positioned letters that have the same intensity in within each shape, (b) randomly positioned letters or (c) circles that have different intensities, or (d) randomly rotated card suits. Furthermore, when accelerated by $R = 4$, even without ACS data (e), these perturbations can be resolved via parallel imaging methods such as CG-SENSE: (f) 20 iterations, (g) 40 iterations, (h) 80 iterations.

for EI (Chen et al., 2021), and employed 3 rotations along with 2 flips. For CUPID, dual-tree complex wavelet transform (DTCWT) which provides an over-complete representation (Selesnick et al., 2005; Cotter, 2020) was selected as the sparsifying transform (\mathbf{W}) in (8). Furthermore, $\mathbf{x}^{(0)}$ in (8), i.e. the initial estimate prior to any reweighting, was calculated using a CS approach as mentioned in Section 3 This was implemented using (11) with 100 iterations, with 30 CG steps for data fidelity and $0.1 \cdot \|\mathbf{W}\mathbf{x}_p\|_\infty$ as the soft thresholding parameter.

B PERTURBATION STRATEGIES

B.1 CHOICE FOR NUMBER OF PERTURBATION PATTERNS

The empirical expectation that approximates the one in (9) is expected to converge to the true expectation as we introduce more perturbation patterns and randomness over the choice of \mathbf{p} . Fig. 7 shows the zero-shot fine-tuning results of CUPID with $K \in \{1, 3, 6, 10\}$, while Fig. 8(a) and Fig. 8(b) illustrates the corresponding PSNR and SSIM curves throughout the training epochs, respectively. As expected, using a single pattern could not capture the true mean and exhibits artifacts. As we introduce more perturbations, we reduce the artifacts and noise amplification. At a certain point, increasing the number of perturbations becomes counterproductive, yielding only marginal gains while significantly increasing the computation time. Thus, we opted to use 6 distinct \mathbf{p}_k patterns throughout our study as it offers the optimal trade-off.

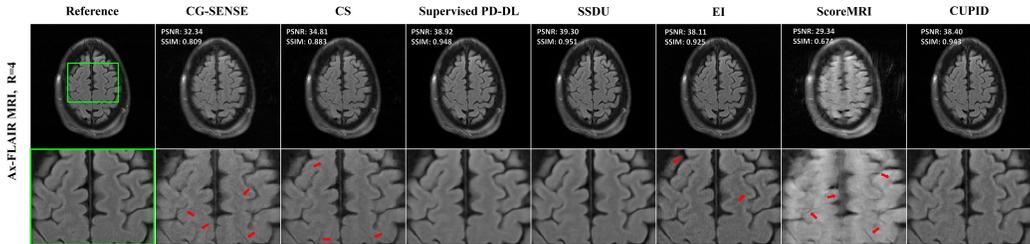


Figure 10: Representative ax-FLAIR database reconstructions based on various learning and reconstruction strategies. CG-SENSE, CS, EI and ScoreMRI exhibit artifacts. On the other hand, CUPID surpasses them with high-fidelity reconstructions, closely matching supervised and self-supervised methods that requires raw k-space data during training.

Table 2: Quantitative results with standard error of the mean included. Each method is tested on Coronal PD, Coronal PD-FS and Ax-FLAIR datasets using equispaced undersampling pattern at $R = 4$.

Method	Coronal PD		Coronal PD-FS		Ax-FLAIR	
	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow
Supervised	40.95 \pm 2.95	0.964 \pm 0.016	35.89 \pm 2.72	0.859 \pm 0.089	36.69 \pm 1.73	0.926 \pm 0.016
SSDU	40.12 \pm 2.95	0.956 \pm 0.019	35.35 \pm 2.73	0.856 \pm 0.091	36.98 \pm 1.83	0.929 \pm 0.016
EI	33.29 \pm 6.65	0.919 \pm 0.034	29.86 \pm 6.59	0.704 \pm 0.156	35.48 \pm 4.55	0.908 \pm 0.028
ScoreMRI	32.84 \pm 4.48	0.812 \pm 0.035	28.18 \pm 4.25	0.684 \pm 0.142	28.17 \pm 3.84	0.774 \pm 0.034
CS	36.71 \pm 2.86	0.917 \pm 0.036	32.30 \pm 2.68	0.749 \pm 0.158	32.93 \pm 2.03	0.865 \pm 0.028
CG-SENSE	35.38 \pm 3.04	0.873 \pm 0.039	30.13 \pm 2.73	0.702 \pm 0.164	30.25 \pm 2.12	0.801 \pm 0.029
CUPID (ours)	39.28 \pm 2.91	0.951 \pm 0.022	34.71 \pm 2.61	0.840 \pm 0.093	36.02 \pm 1.70	0.920 \pm 0.018

B.2 DESIGN ALTERNATIVES FOR PERTURBATIONS

As stated in Sec. 3, added perturbations may consist of several different structures. Fig. 9 provides some of these perturbation examples, [an illustration of how the perturbation looks with undersampling, and how they are recovered perfectly through conventional parallel imaging methods](#). We note that there was no task-specific perturbation that we used, meaning that the perturbations selected from the same set were applied to all datasets given that the created perturbations do not create fold-overs at R-fold which result in artifacts. Note the latter condition means they should be recoverable through parallel imaging reconstruction. [Finally, we note that when calculating the sample mean estimate for \(9\), intensity of the perturbations was empirically found to be more important than their shapes/orientations. Specifically, we observed that varying it randomly within the perturbation, as in Fig. 9b-d, leads to improved reconstruction outcomes.](#)

C MORE RESULTS FROM THE RETROSPECTIVE STUDY

We further include the representative reconstructions from the Ax-FLAIR dataset in Fig. 10.

D FURTHER QUANTITATIVE RESULTS

[A more detailed version of Tab. 1, incorporating the standard error of the mean, is provided in Tab. 2.](#)

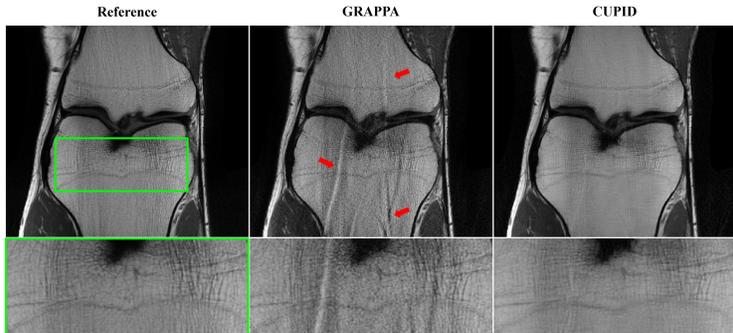


Figure 11: Representative reconstructions for CUPID with x_{PI} reconstructed using GRAPPA on coronal PD knee MRI using $R = 4$ uniform undersampling. GRAPPA exhibits aliasing and noise artifacts at this high acceleration rate. PD-DL network trained with a CUPID implementation that only has access to this GRAPPA reconstruction improves on it, reducing these artifacts. This highlights the compatibility of CUPID with different parallel imaging reconstructions.

918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

E COMPATIBILITY WITH VARIOUS PARALLEL IMAGING RECONSTRUCTIONS

Vendor reconstructions typically use different parallel imaging techniques. For our retrospective studies, we used CG-SENSE (or equivalently SENSE) (Pruessmann et al., 1999) because it naturally fits with the DF units in the unrolled network, and it is commonly used in clinical settings, alongside GRAPPA (Griswold et al., 2002). However, we emphasize that our method does not make assumptions about the specific reconstruction method used by the vendor; instead, it assumes that parallel imaging can resolve the perturbations, which is ensured by designing them in a manner that prevents fold-over aliasing artifacts from overlapping.

To further validate this, we include representative CUPID reconstruction results in Fig. 11 where x_{PI} is generated via GRAPPA (Griswold et al., 2002), demonstrating that CUPID is compatible with different types of parallel imaging reconstructions as input. We further note that the prospective study also used GRAPPA reconstruction as input, as this is the reconstruction provided by the vendor used in our institution.