

Validation and Generalizability of Self-Supervised Image Reconstruction Methods for Undersampled MRI

Thomas Yu^{1,2}, Tom Hilbert^{1,3,4} *, Gian Franco Piredda^{1,3,4}, Arun Joseph^{5,6,7}, Gabriele Bonanno^{5,6,7}, Salim Zenkhri⁴, Patrick Omoumi⁴, Meritxell Bach Cuadra^{2,4,1}, Erick Jorge Canales-Rodríguez¹, Tobias Kober^{1,3,4}, Jean-Philippe Thiran^{1,2,4}

1. Signal Processing Laboratory 5 (LTS5), École Polytechnique Fédérale de Lausanne, Lausanne (EPFL), Switzerland
2. CIBM, Center for Biomedical Imaging, Switzerland
3. Advanced Clinical Imaging Technology, Siemens Healthineers International AG, Lausanne, Switzerland
4. Department of Radiology, Lausanne University Hospital and University of Lausanne, Switzerland
5. Advanced Clinical Imaging Technology, Siemens Healthineers International AG, Bern, Switzerland
6. Translational Imaging Center, sitem-insel, Bern, Switzerland
7. Magnetic Resonance Methodology, Institute of Diagnostic and Interventional Neuro-radiology, University of Bern, Bern, Switzerland

Abstract

Deep learning methods have become the state of the art for undersampled MR reconstruction. Particularly for cases where it is infeasible or impossible for ground truth, fully sampled data to be acquired, self-supervised machine learning methods for reconstruction are becoming increasingly used. However potential issues in the validation of such methods, as well as their generalizability, remain underexplored. In this paper, we investigate important aspects of the validation of self-supervised algorithms for reconstruction of undersampled MR images: quantitative evaluation of prospective reconstructions, potential differences between prospective and retrospective reconstructions, suitability of commonly used quantitative metrics, and generalizability. Two self-supervised algorithms based on self-supervised denoising and the deep image prior were investigated. These methods are compared to a least squares fitting and a compressed sensing reconstruction using in-vivo and phantom data. Their generalizability was tested with prospectively under-sampled data from experimental conditions different to the training. We show that prospective reconstructions can exhibit significant distortion relative to retrospective reconstructions/ground truth. Furthermore, pixel-wise quantitative metrics may not capture differences in perceptual quality accurately, in contrast to a perceptual metric. In addition, all methods

*. Corresponding Author

showed potential for generalization; however, generalizability is more affected by changes in anatomy/contrast than other changes. We further showed that no-reference image metrics correspond well with human rating of image quality for studying generalizability. Finally, we showed that a well-tuned compressed sensing reconstruction and learned denoising perform similarly on all data. The datasets acquired for this paper will be made available online; see <https://www.melba-journal.org/papers/2022:022.html> for details.

Keywords: Deep Learning, Self-Supervised Learning, MR Image Reconstruction, Validation, Generalizability

1. Introduction

Since the introduction of MRI, methods for image reconstruction have evolved with acquisition acceleration and have seen great advances with parallel imaging techniques such as sensitivity encoding (SENSE) (Pruessmann et al., 1999) and generalized auto-calibrating partially parallel acquisition (GRAPPA) (Griswold et al., 2002). While parallel imaging reliably accelerates clinical contrasts by factors of two to three, more recent methods such as compressed sensing (CS) have achieved even higher acceleration factors (Lustig et al., 2007). Now, supervised deep learning methods reign as the state of the art in the reconstruction of accelerated acquisitions (Knoll et al., 2020a; Hammernik and Knoll, 2020; Sun et al., 2016). However, these supervised methods require a non-trivial amount of fully sampled data to use as ground truth/target, which can be difficult or infeasible to obtain depending on the type of acquisition. Consequently, there has been interest in unsupervised or self-supervised, deep learning approaches which train solely on accelerated acquisitions, with no need for ground truth, fully sampled data (Liu et al., 2020; Yaman et al., 2020; Heckel and Hand, 2019; Akçakaya et al., 2021).

However, the validation of these methods is generally done by quantitative evaluation through pixel-wise metrics on **retrospectively undersampled** acquisitions (i.e., artificial undersampling of a fully sampled dataset), sometimes accompanied by qualitative evaluation on datasets where no ground truth is available. This limitation may stem from commonly used datasets (Epperson et al., 2013; Knoll et al., 2020b) being fully sampled, as well as difficulties in acquiring datasets which contain both fully sampled and prospectively accelerated scans without motion corruption. However, this neglects quantitative evaluation of reconstructions from **prospectively undersampled** data, the clinically relevant scenario, as well as potential differences between prospective and retrospective reconstructions; furthermore, the pixel-wise metrics generally used may not correlate well with the perceptual quality of the images. This point is crucial for clinical deployment as even if different methods can be robustly ranked using retrospective data, the image quality from prospective data from the different methods may be unsuitable for clinical use. Furthermore, if these techniques will be used in future clinical routines, they likely will be subject to variations of data quality and content. For example, different surface coils, parameter differences between centers or even the use of the same sequence on different organs. Therefore, the generalizability, i.e., inference data different from the training/tuning data (e.g. in terms of field strength, sequence parameters, motion, anatomy, etc.), using prospective data is of interest, both for investigating robustness and for testing the limits of self-supervised methods. Furthermore, while prospective reconstructions are generally evaluated using qualitative

rating, we evaluated the potential for using no-reference image metrics for a quantitative evaluation.

1.1 Contributions

In this work, we fixed an MR sequence of interest for which extensive, clinical acquisition of fully sampled data is infeasible and conducted an extensive, realistic validation of state of the art self-supervised reconstruction methods through two novel, overarching experiments.

1. In contrast to the literature, we acquired phantom data with both full sampling and prospective acceleration. This allowed us to quantitatively and qualitatively evaluate both prospective and retrospective reconstructions using both pixel-wise and perceptual metrics for fidelity to ground truth, allowing us to study them individually as well as to see any relevant differences.
2. In contrast to the literature, we tested the generalizability of the methods using an extensive, prospectively accelerated dataset with changes in contrast, hardware, field strength, and anatomy. Furthermore, we evaluated the results both quantitatively, using no-reference image quality metrics, and qualitatively, using rating by MR scientists and a radiologist.

2. Theory

The self-supervised, machine-learning based methods we examine in this paper rely on two powerful ideas drawn from machine learning: self-supervised denoising and restriction to the range of convolutional neural networks (CNN) as an effective prior for image reconstruction. We chose these methods for validation as these ideas have been shown to be both empirically effective and theoretically well founded, making them attractive for clinical use. In Figure 1, we show an overview of the different methods used in this paper. We begin with the basic inverse problem formulation of MR image reconstruction. Let $\mathbf{y}_i, \mathbf{n}_i$ denote the undersampled MR measurements and Gaussian noise respectively, from the i th coil element and \mathbf{x} denote the underlying image. These quantities are related by:

$$\mathbf{y}_i = A_i \mathbf{x} + \mathbf{n}_i, \quad (1)$$

$$A_i = M \circ F \circ S_i \quad (2)$$

where M is the element-wise multiplication by a mask (corresponding to the location of the undersampled measurements), F denotes the Fourier transform, and S_i denotes element-wise multiplication by the i th sensitivity map. The classical regularized reconstruction of \mathbf{x} is the solution of an optimization problem

$$\mathbf{x} = \arg \min_{\mathbf{x}'} D(\mathbf{x}', \mathbf{y}) + \lambda R(\mathbf{x}'), \quad (3)$$

where $D(\mathbf{x}, \mathbf{y})$ measures the consistency of the solution to the data (e.g. $\|A_i \mathbf{x} - \mathbf{y}_i\|^2$), $R(\mathbf{x})$ is a regularization function, which, for example, prevents overfitting to the noise, and λ is the regularization parameter. In combination with incoherently undersampled measurements, compressed sensing reconstructions have been shown to effectively reconstruct the

underlying images by setting $R(\mathbf{x})$ to encourage sparsity of \mathbf{x} in a set domain (Lustig et al., 2007). Many state of the art deep learning methods, both supervised and unsupervised, implicitly or explicitly parametrize $R(\mathbf{x})$ with a neural network. In this work, we choose to compare two state-of-the-art self-supervised approaches which operate by orthogonal, well-founded theoretical principles with impressive empirical performance.

2.1 DeepDecoder

The first self-supervised method we examine is called DeepDecoder. DeepDecoder is based on a seminal work in the machine learning literature called Deep Image Prior (DIP) (Ulyanov et al., 2018) which showed that untrained CNNs could be used to effectively solve inverse problems without ground truth. Concretely, let f_θ denote a randomly initialized CNN with parameters θ . Let \mathbf{z} be a sample of a random, Gaussian vector. Then DIP solves Equation 3 by

$$\mathbf{x} = f_\theta(\mathbf{z}) = \arg \min_{\theta'} \|A_i f_{\theta'}(\mathbf{z}) - \mathbf{y}_i\|^2 \quad (4)$$

This formulation is equivalent to setting $R(\mathbf{x})$ to the indicator function with support over the range of the neural network; this assumes that the convolutional network f_θ itself provides a strong prior on the space of image solutions, such that only the data consistency term needs to be minimized. However, since only the noisy signal \mathbf{y} is used during training, minimization can overfit the noise in the signal, depending on the inverse problem being solved (e.g. denoising, super-resolution), thus requiring early stopping (Ulyanov et al., 2018). DeepDecoder (Heckel and Hand, 2019) is a CNN with a simplified architecture (only upsampling units, pixel-wise linear combination of channels, ReLU activation, and channel-wise normalization) which is amenable to theoretical analysis and was shown to be competitive with other architectures for solving inverse problems in a DIP framework.

In (Heckel and Soltanolkotabi, 2020), the authors theoretically showed that for the case of image recovery from compressed sensing measurements, CNNs (in particular, CNNs with the structure of DeepDecoder) are self-regularizing with respect to noise and can simply be trained to convergence with gradient descent without early stopping or additional regularization, provided that the true, underlying image has sufficient smoothness/structure. In a knee MR example, they showed that early stopping would have only provided a marginally better solution than running to convergence. Hence, from a theoretical and practical standpoint, DeepDecoder is attractive for self-supervised reconstruction from undersampled measurements. We emphasize that DeepDecoder entails training a separate network for each separate acquisition/slice, rather than training a single network over a dataset of under-sampled acquisitions.

2.2 Self-supervised learning via data under-sampling

The second self-supervised method we examine is called Self-supervised learning via data under-sampling (SSDU). SSDU uses an unrolled, iterative architecture, with alternating neural network and data consistency modules, to reconstruct MR images using only under-sampled measurements, with the adjoint image corresponding to the input k-space measurements as an initial guess. It solves Eqn 3 using an iterative, variable splitting approach

where the k th iteration consists of

$$\hat{\mathbf{x}}^k = \text{CNN}(\mathbf{x}^{k-1}) \quad (5)$$

$$\mathbf{x}^k = \arg \min_{\mathbf{x}'} \|A_i \mathbf{x}' - \mathbf{y}_i\|^2 + \lambda \|\mathbf{x}' - \hat{\mathbf{x}}^k\|^2. \quad (6)$$

where the superscript denotes the iteration, CNN denotes a generic CNN, and $\hat{\mathbf{x}}^k$ denotes an auxiliary variable. The regularization parameter λ is learned during training. Let f_{SSDU} denote the function defined by the unrolled network. In each training step of SSDU, the k-space of the data is split into two, random disjoint sets, denoted by \mathbf{y}_Θ and \mathbf{y}_Λ . \mathbf{y}_Θ is passed to the unrolled network as input. The loss function for SSDU compares the simulated k-space measurements of the corresponding image output $f_{SSDU}(\mathbf{y}_\Theta)$ to \mathbf{y}_Λ :

$$L(\mathbf{y}_\Lambda, A_\Lambda f_{SSDU}(\mathbf{y}_\Theta)) \quad (7)$$

where A_Λ is the measurement operator corresponding to sampling the locations of Λ , and L is an equally weighted combination of the L_1 and L_2 loss. Hence, during each training step, f_{SSDU} only sees information from \mathbf{y}_Θ , and the loss is only computed over a disjoint set \mathbf{y}_Λ . We note that at inference time, the entire, acquired k-space measurements are given as input. While the authors of SSDU give an intuitive explanation of this approach as similar to cross validation in order to prevent overfitting to noise or learning the identity, results from the machine learning literature on blind, signal denoising can help give a theoretical explanation.

In the Noise2Self framework (Batson and Royer, 2019), the authors prove that a neural network can be trained to denoise a noisy signal, using solely the noisy signal for training. In the following, we describe a special case of the general theory proven in (Batson and Royer, 2019). Let $\mathbf{y}^\delta = \mathbf{y} + \mathbf{n}$ denote a noisy signal, where \mathbf{y} , \mathbf{n} are the noise-free signal and Gaussian noise respectively. Partition \mathbf{y}^δ into disjoint sets, \mathbf{y}_J^δ and $\mathbf{y}_{J^C}^\delta$, where the subscript indicates restriction of the corresponding vectors to the disjoint subsets of indices J, J^C , with other indices being zero-filled. Then the authors showed that that a neural network (denoted as f) can be trained to denoise the noisy signal, using solely the noisy signal, by using the following loss function:

$$L(f) = \sum_J \mathbb{E} \|f_J(\mathbf{y}_{J^c}^\delta) - \mathbf{y}_J^\delta\|^2 = \sum_J \mathbb{E} \|f_J(\mathbf{y}_{J^c}^\delta) - \mathbf{y}_J\|^2 + \|\mathbf{y}^\delta - \mathbf{y}\|^2 \quad (8)$$

We emphasize that the right hand side of Equation 8 is composed of the mean squared error between the signal predicted by the network and the ground-truth signal and a constant independent of the network. **Hence the Noise2Self strategy allows to minimize the error between the predicted signal and the ground truth signal with only access to the noisy signal, by iteratively giving a partition of the noisy signal as input to f and computing the MSE over a disjoint partition.** Identifying Θ, Λ with J, J^c , we can see that the training of SSDU conforms to the Noise2Self framework with the k-space measurements acting as the noisy signal, albeit with SSDU using an L_1 loss in addition to the L_2 loss. Thus, SSDU takes as input the noisy, acquired k-space measurements, and is optimized to output an image whose simulated k-space measurements are the acquired k-space measurements **without noise**. In this way, SSDU avoids overfitting to noise.

This, combined with the powerful image prior from using a CNN as the neural network as well as the interleaving of the data consistency term, explains SSDU’s demonstrated ability to provide denoised images which retain image sharpness, as compared to traditional methods. We can interpret SSDU as an iterative method which interleaves the application of a denoising network and a data consistency step. We note in contrast to DeepDecoder, that we can train different networks for separate acquisitions or train a single, reusable network on a dataset of undersampled acquisitions. In this paper, we do the latter.

In conclusion, both self-supervised approaches accomplish noise robust MR reconstruction using only noisy, undersampled MR measurements;

3. Methods

In the following experiments, we compare four image reconstruction methods:

1. **CG-SENSE**, which solves Equation 3 with no regularization using the conjugate gradient algorithm; this is a least squares fit to the acquired data similar to the description in (Pruessmann et al., 2001).
2. **CS-L1Wavelet**, where we solve Equation 3 with a compressed sensing reconstruction, with $R(\mathbf{x}) = \|W\mathbf{x}\|_1$, where W is a wavelet transform operator. We set the regularization parameter λ to 2.3e-4 according to a Noise2Self tuning described in the appendix.
3. **DeepDecoder** with a depth/width of 300/10 and Gaussian input of size (10,10).
4. **SSDU**, where we use a U-Net (Ronneberger et al., 2015) with 12 channels and 4 downsampling/upsampling layers. Training (Θ) and testing (Λ) masks are randomly sampled uniformly, with a split of 60 and 40 percent respectively.

We used Sigpy(Ong and Lustig, 2019) for the computation of CS-L1Wavelet and ES-PiRiT(Uecker et al., 2014) sensitivity maps. We implemented CG-SENSE and SSDU in Pytorch (Paszke et al., 2019), and we used Github implementations of DeepDecoder[†] and U-Net[‡]. We used Adam (Kingma and Ba, 2014) to optimize both SSDU and DeepDecoder. SSDU was trained until convergence (10 epochs) with a learning rate of 0.5e-4. For each subject, DeepDecoder was optimized using the acceleration strategy in (Darestani and Heckel, 2021); a single slice for each subject is optimized to convergence (over 10,000 iterations) from a random initialization. All other slices are optimized for 1,000 iterations, initialized with the network model from this single slice. All training and inference was done on a NVIDIA Quadro RTX 8000 with 45GB of RAM.

3.1 Training Data and Hyperparameter Tuning

To mimic a realistic scenario with a sequence for which fully sampled, ground truth data is difficult/infeasible to acquire, and where the training dataset is limited in size and variability, we acquired for ten healthy subjects a 5x accelerated 3D MPRAGE prototype sequence

[†]. https://github.com/MLI-lab/cs_deep_decoder

[‡]. <https://github.com/facebookresearch/fastMRI>

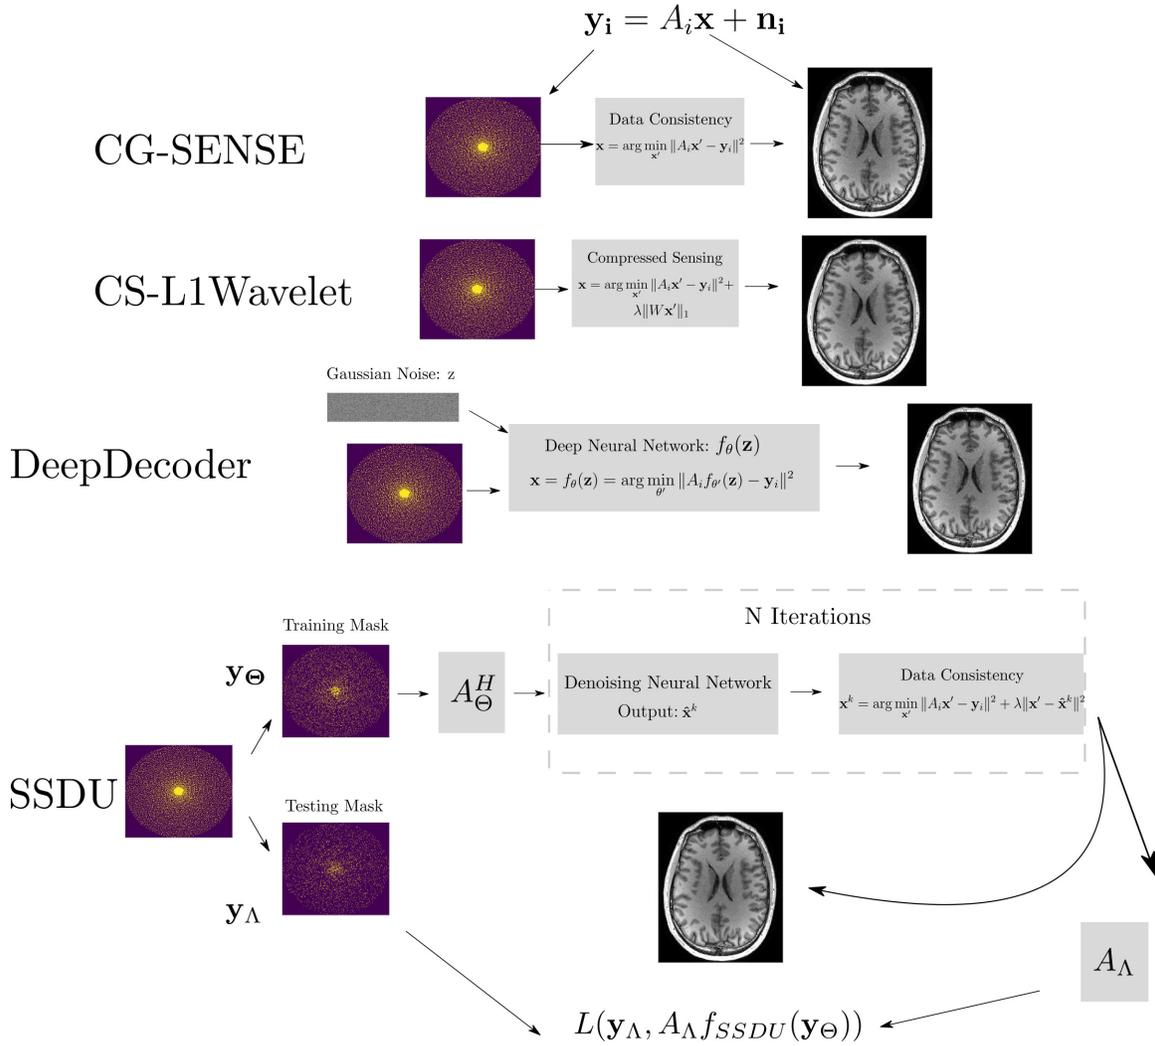


Figure 1: An overview of the basic formulation of the MR reconstruction inverse problem, as well as how each method in the paper solves the inverse problem.

(Mussard et al., 2020) of the brain at 3T (MAGNETOM Prisma^{Fit}, Siemens Healthcare, Erlangen, Germany) using a 64ch Rx Head/Neck coil. These incoherently undersampled data were used for training/tuning the hyperparameters of all reconstruction methods. In what follows, all training/inference is done on 2D slices of both phase-encoding directions formed from performing the inverse Fourier transform along the readout direction. **We emphasize that in the absence of prior knowledge/heuristics, the hyperparameters of the methods should also be tuned in a self-supervised way, as the common method for hyperparameter tuning, i.e. using a hold-out set of data for which the ground truth is known, is not available in our scenario.** We use the Noise2Self framework, which also underlies SSDU, for selecting hyperparameters (regularization parameter of CS-L1Wavelet and the network parameters of DeepDecoder and SSDU), as it optimizes for preventing overfitting to the noise in the measurements. Details on the hyperparameter tuning can be found in the Appendix.

3.2 Validation using Prospectively Accelerated and Fully Sampled Data

In our first experiment, using the aforementioned 3D MPRAGE prototype sequence used for acquiring the training/tuning data, we acquired both fully sampled and 5x prospectively accelerated scans of the following:

1. Siemens multi-purpose phantom E-38-19-195-K2130 filled with $MnCl_2 \cdot 4H_2O$ doped water
2. Assortment of fruits/vegetables (Pineapple, tomatoes, onions, brussel sprouts)

This allowed us to reconstruct prospective, retrospective (applying the same mask as in prospective sampling on the fully sampled data), and fully-sampled images.

No in-vivo data was used in this experiment since subject motion could bias the results. Furthermore, we used fruits/vegetables as a second phantom since they have more complex structures than a water filled container.

3.2.1 QUANTITATIVE ASSESSMENT

First, we qualitatively compared the results through visual inspection. Second, we quantitatively compare reconstructions to the ground truth using Peak Signal to Noise Ratio (PSNR) (Salomon, 2004), the Structural Similarity Index Measure (SSIM) (Wang et al., 2004), and a metric we will call the Perceptual Distance (PercDis) score. While the first two are commonly used metrics in MR image reconstruction/image reconstruction in general, the PercDis score comes from computer vision (super-resolution, style transfer, etc), where it is called the perceptual loss (Johnson et al., 2016); the distance between two images is defined as the L_1 distance between the respective induced features from intermediate layers of a pretrained image classification network. The scores of center cropped slices, along the read-out direction, are averaged for the final score.

3.3 Generalizability of Self-Supervised Reconstruction Methods

In our second experiment, we examined the generalizability of the reconstruction methods. To that end, we scanned three, healthy subjects with the following prospectively accelerated sequences(anatomy):

1. 1.5T MPRAGE (Brain)
2. 3T MPRAGE (Brain)
3. 7T MPRAGE (Brain)
4. 3T MPRAGE with 1Tx/20Rx Coil (Brain)
5. 3T MPRAGE with Subject Motion (Brain)
6. 3T MPRAGE with Different Parameters (Brain)
7. 3T, T_1 SPACE (Brain)
8. 3T, T_2 FLAIR SPACE (Brain)
9. 3T, PD SPACE (Knee)
10. 3T, T_2 SPACE (Knee)

The brain scans at 1.5T, 3T and 7T (MAGNETOM Sola, Vida, and Terra, Siemens Healthcare, Erlangen, Germany) were done using a 1Tx/20Rx, 1Tx/64Rx (unless otherwise stated), and 8pTx/32Rx (Nova Medical, Wilmington, MA, USA) head coil, respectively. The knee scans at 3T were done with a 1Tx/18Rx coil. All detailed sequence parameters can be found in the Appendix in Table 2.

As ground truth data is not available since motion would render quantitative comparison difficult due to blurring from image co-registration, we evaluated the reconstructions from the above data quantitatively through no-reference image quality metrics and qualitatively through rating by four MR scientists and a radiologist. In total, 120 reconstructions (40 per subject) were evaluated.

3.3.1 NO-REFERENCE IMAGE METRICS

No-reference image quality metrics quantify the quality of a given image (i.e. blurriness, noise) using only its statistical features in a way that correlates with the perceptual quality of a human observer. They have been shown to potentially be useful for MR/medical image evaluation without ground truth (Woodard and Carley-Spencer, 2006; Zhang et al., 2018); we use the following three metrics: a metric used originally for assessing the quality of JPEG-compressed images which we call NRJPEG (Wang et al., 2002), the Blind/Referenceless Image Spatial Quality Evaluator (BRISQUE) (Mittal et al., 2011), and Perception based Image Quality Evaluator (PIQE) (Venkatanath et al., 2015). BRISQUE and PIQE have also been used in other image reconstruction challenges where the ground truth is not available, such as super-resolution (Lugmayr et al., 2020). The metrics were calculated for the central 100 slices (along the read-out direction) of each reconstruction.

3.3.2 HUMAN QUALITY RATING

The human quality rating was done according to (Hammernik et al., 2018) by four experienced MR scientists and a radiologist. Using a 4-point ordinal scale, reconstructed images were evaluated for sharpness (1: no blurring, 2: mild blurring, 3: moderate blurring, 4: severe blurring), SNR (1: excellent, 2: good, 3: fair, 4: poor), presence of aliasing artifacts (1: none, 2: mild, 3: moderate, 4: severe) and overall image quality (1: excellent, 2: good, 3: fair, 4: poor). Raters were blinded to the reconstruction method.

3.4 Statistical Significance

For all quantitative metrics/ratings, we use the Wilcoxon signed rank test with significance level $\frac{0.05}{6}$ (Bonferroni correction with 6 pair-wise comparisons among the 4 methods) to determine statistical significance.

4. Results

In general, perceptually, CG-SENSE produces noisy but sharp images since it is not regularized. DeepDecoder produces smoother reconstructions with spatially varying noise behavior and sharpness, e.g Figure 2 (yellow arrows). CS-L1Wavelet and SSDU produce similar images, smoother than those of CG-SENSE with comparable sharpness; however, CS-L1Wavelet exhibits more artifacts, e.g Figure 2 (red arrows).

4.1 Validation Using Prospectively Accelerated and Fully Sampled Data

In Fig. 2 and Fig. 3, we can see spatial distortions of hyper/hypo-intense features in the prospective reconstructions and changes in contrast in comparison to the ground truth reconstruction; this distortion is not present in the retrospective reconstructions; however, they are similar across all reconstruction methods.

Retrospective reconstructions have significantly higher mean scores for all metrics in comparison to the prospective reconstructions in both acquisitions (see Table 1).

Comparing the methods, in the phantom, the prospective/retrospective reconstructions of DeepDecoder have the highest pixel-wise fidelity to the ground truth with a mean PSNR of (18.67/23.44) and SSIM of (0.49/0.52); however, qualitatively, it has more spatially varying oversmoothing than those of CS-L1Wavelet and SSDU. SSDU and CS-L1Wavelet perform similarly, with the highest qualitative similarity to the ground truth, with SSDU having a higher mean PSNR overall (17.79/21.95). In contrast to the PSNR/SSIM results, with the PercDis score, SSDU has the highest fidelity to the ground truth (0.63/0.61).

Qualitatively and quantitatively (with PSNR and SSIM), the differences between the methods are much less in the fruits/vegetables. The main qualitative difference is the greater denoising capabilities of SSDU and CS-L1Wavelet in comparison to CG-SENSE and DeepDecoder. Quantitatively, there are only minor differences between the methods with respect to PSNR and SSIM. In contrast, the PercDis scores clearly indicate that CS-L1Wavelet and SSDU (with similar scores) are perceptually more similar to the ground truth than CG-SENSE and DeepDecoder (with similar scores).

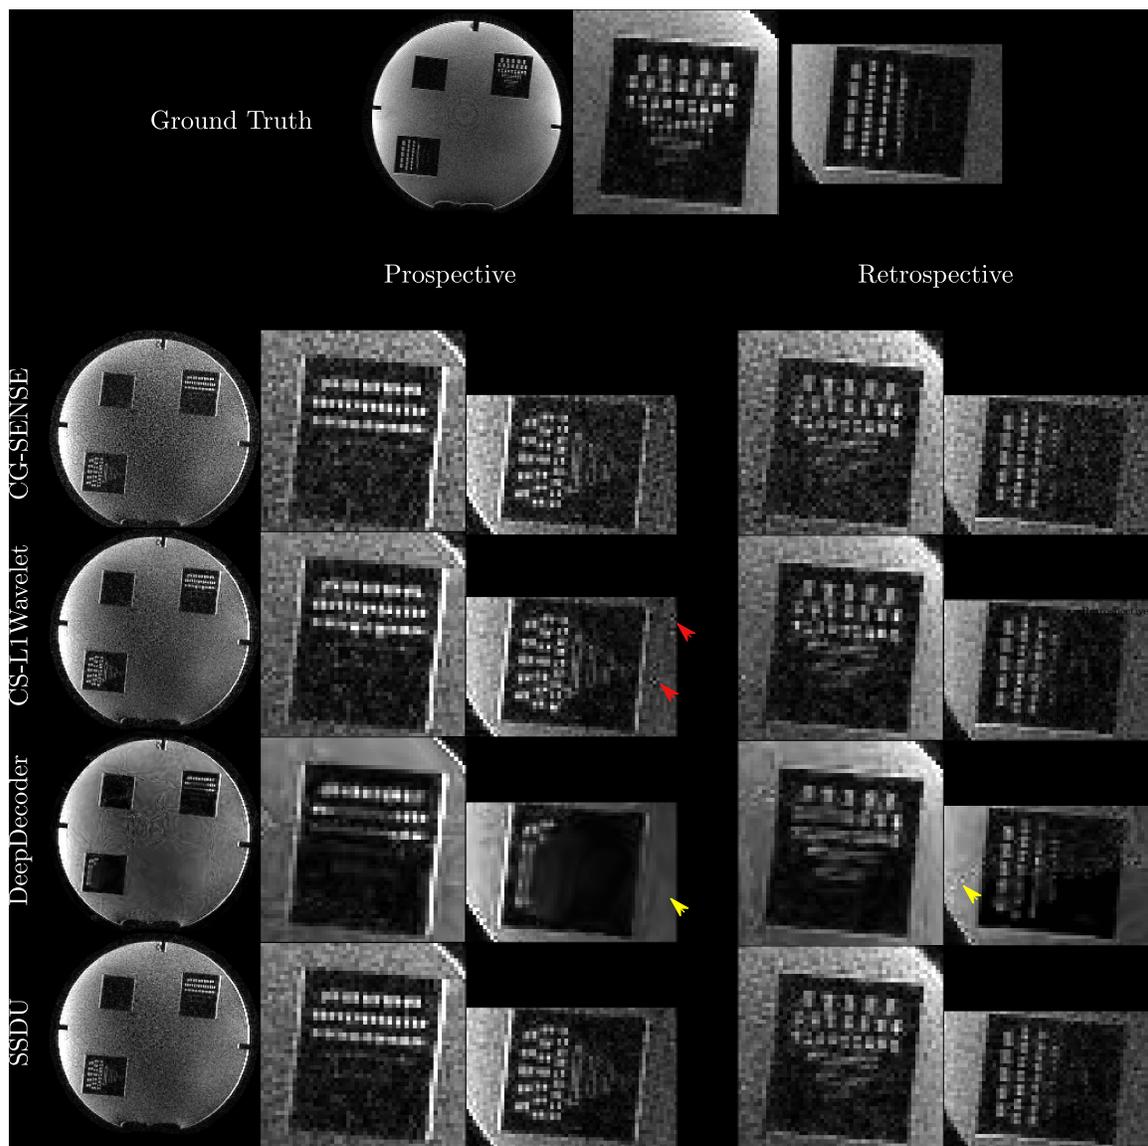


Figure 2: Ground truth images and reconstructed images using prospectively and retrospectively accelerated data from the multi-purpose phantom, scanned with a MPRAGE sequence at 3T. Reconstructions from prospectively accelerated data are distorted (see closeups) relative to the ground truth/retrospective reconstructions. DeepDecoder exhibits spatially varying smoothness/distortion (see yellow arrows) relative to CS-L1Wavelet and SSDU which have similar scores/appearance, although CS-L1Wavelet has more artifacts (see red arrows). CG-SENSE produces noisy but sharp reconstructions, while CS-L1Wavelet and SSDU reduce noise but preserve sharpness relative to CG-SENSE.

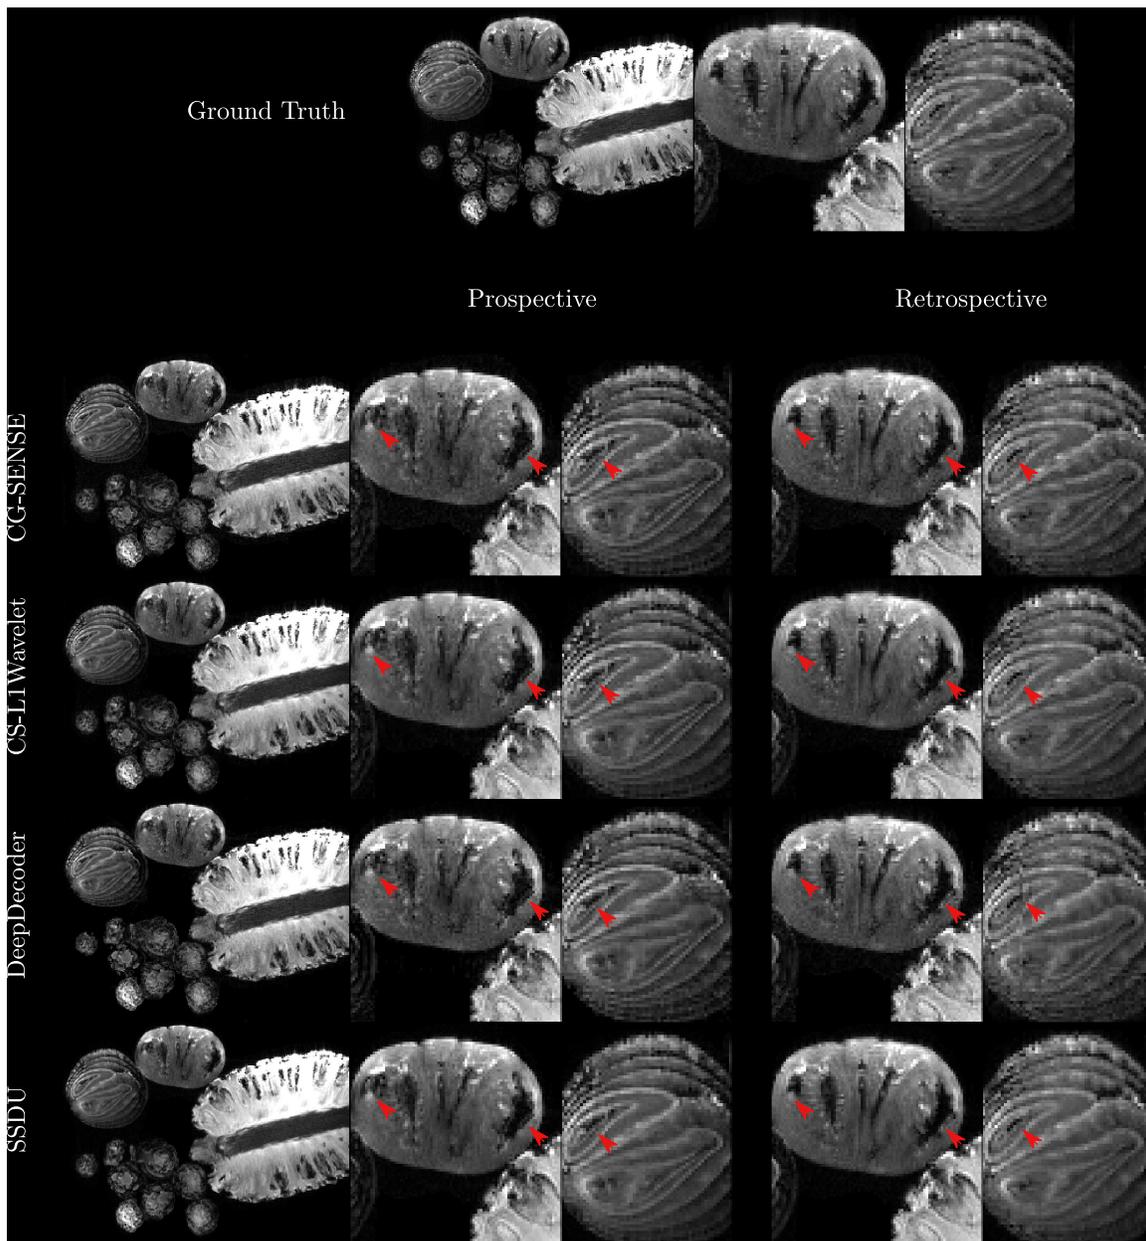


Figure 3: Ground truth images as well as reconstructed images using prospectively and retrospectively accelerated data from the fruits/vegetables, scanned with a MPRAGE sequence at 3T. Reconstructions from prospectively accelerated data are distorted in hypointense regions (see closeup/red arrows) relative to the ground truth/retrospective reconstructions. Qualitatively, the main difference between the methods are between CS-L1Wavelet/SSDU and CG-SENSE/DeepDecoder, where the former group is smoother than the latter.

PSNR \uparrow	Phantom		Fruits/Vegetables	
(μ, σ)	Prospective	Retrospective	Prospective	Retrospective
CG-SENSE	(13.54,9.69)	(16.82,12.17)	(33.4,4.86)	(39.3,5.73)
CS-L1Wavelet	(16.0,8.36)	(20.62,11.65)	(33.59,3.83)	(38.88,4.44)
DeepDecoder	(18.67,6.65)	(23.44,10.66)	(33.65,2.86)	(38.25,4.42)
SSDU	(17.79,6.9)	(21.95,10.09)	(33.88,3.75)	(39.49,4.27)

SSIM \uparrow	Phantom		Fruits/Vegetables	
(μ, σ)	Prospective	Retrospective	Prospective	Retrospective
CG-SENSE	(0.35,0.18)	(0.42,0.23)	(0.92,0.09)	(0.95,0.09)
CS-L1Wavelet	(0.4,0.21)	(0.47,0.27)	(0.93,0.08)	(0.96,0.08)
DeepDecoder	(0.49,0.22)	(0.52,0.28)	(0.93,0.04)	(0.95,0.08)
SSDU	(0.41,0.22)	(0.47,0.27)	(0.93,0.08)	(0.96,0.08)

PercDis \downarrow	Phantom		Fruits/Vegetables	
(μ, σ)	Prospective	Retrospective	Prospective	Retrospective
CG-SENSE	(1.05,0.08)	(1.02,0.06)	(0.45,0.16)	(0.29,0.09)
CS-L1Wavelet	(0.84,0.08)	(0.79,0.05)	(0.41,0.17)	(0.25,0.09)
DeepDecoder	(0.68,0.15)	(0.64,0.09)	(0.44,0.19)	(0.3,0.11)
SSDU	(0.63,0.13)	(0.61,0.1)	(0.42,0.16)	(0.26,0.09)

Table 1: Mean and standard deviation of PSNR/SSIM/PercDis scores of the reconstructions with respect to the ground truth for the phantom and the fruit/vegetables; arrows beside each metric denote whether higher or lower values are better. PSNR/SSIM/PercDis were calculated over all the slices in the read-out direction with center cropping. We found statistically significant differences between each method for each metric **other than** (CS-L1Wavelet vs SSDU, Retrospective SSIM, Phantom) and (CG-SENSE vs SSDU, Retrospective PSNR, Fruits/Vegetables) Note that while with respect to PSNR/SSIM, DeepDecoder performs the best in the phantom, and all methods perform similarly in Fruits/Vegetables. In contrast, with respect to the PercDis score, SSDU performs the best in both cases by larger relative margins than with PSNR/SSIM.

4.2 Generalizability

Figures 4, 5 show axial MPRAGE brain slices at the different field strengths and corresponding closeups of the cerebellum and the left frontal lobe. Figure 6 shows a sagittal PD knee slice (3T) with closeups of articular cartilage interfaces in sagittal (femur) and axial (patella) views. These show the generalizability of the methods to different magnetic field strengths as well as changes in anatomy and contrast. Example reconstructions for the other sequences can be found in the Appendix 8, 9.

4.2.1 PERCEPTUAL EVALUATION

Qualitatively, we can see from Figures 4, 5, 6 that all methods are able to generalize well (in the sense of approximately preserving performance/appearance on dataset used for training/tuning) to changing field strengths, anatomy, and contrast, although changing anatomy clearly worsened absolute image quality as compared to changing field strength. DeepDecoder preserves its spatially varying smoothing/artifacts, and SSDU/CS-L1Wavelet are able to produce images with less noise and comparable sharpness to CG-SENSE, although CS-L1Wavelet exhibits more artifacts. As expected, the perceptual quality of all methods increase with increasing field strength due to higher spatial resolution. Differences between the methods are less pronounced in the knee scan although overall image quality is worse.

4.2.2 NO-REFERENCE IMAGE QUALITY METRICS

In the first row of Figure 7, we show a bar plot of the scores for the no-reference image quality metrics averaged over all sequences and subjects. In general, CS-L1Wavelet and SSDU have the highest (by a small margin) mean NRJPEG score (10.54/10.39) and lowest, mean BRISQUE (29.35/28.06) and PIQE (25.56/22.87) scores, indicating better image quality in comparison to CG-SENSE and DeepDecoder.

4.2.3 HUMAN RATINGS

In the second row of Figure 7, we show bar plots of the scores from the MR scientists and the radiologist; we pooled the scores of the MR scientists. We see that MR scientists and the radiologist generally agree for evaluating SNR, aliasing, and overall quality, rating CS-L1Wavelet/SSDU as being better than or the same as CG-SENSE/DeepDecoder. We recall that lower ratings correspond to better quality. MR scientists rated CS-L1Wavelet/SSDU with a mean overall quality of (2.09/1.97) as compared to CG-SENSE/DeepDecoder with (2.96/3.57). The radiologist rated CS-L1Wavelet/SSDU with a mean overall quality of (2.73/2.23) as compared to CG-SENSE/DeepDecoder with (3.63/3.87). We note that for both sets of raters, the difference between CS-L1Wavelet and SSDU in overall image quality was found to not be statistically significant. Furthermore, when we restrict our analysis to the average score change between the subgroup of changes in field strength vs. the subgroup of PD Knee/ T_2 Knee scans, the overall image quality rating of CG-SENSE/CS-L1Wavelet/DeepDecoder/SSDU all worsen in the knee scans for the MR scientists, with increases of 0.26, 0.40, 0.11, and 0.79 respectively. In contrast, for the radiologist, this shift results in changes of -0.33, 0.33, -0.16, and 0.83 respectively, indicating that only CS-L1Wavelet and SSDU worsened.

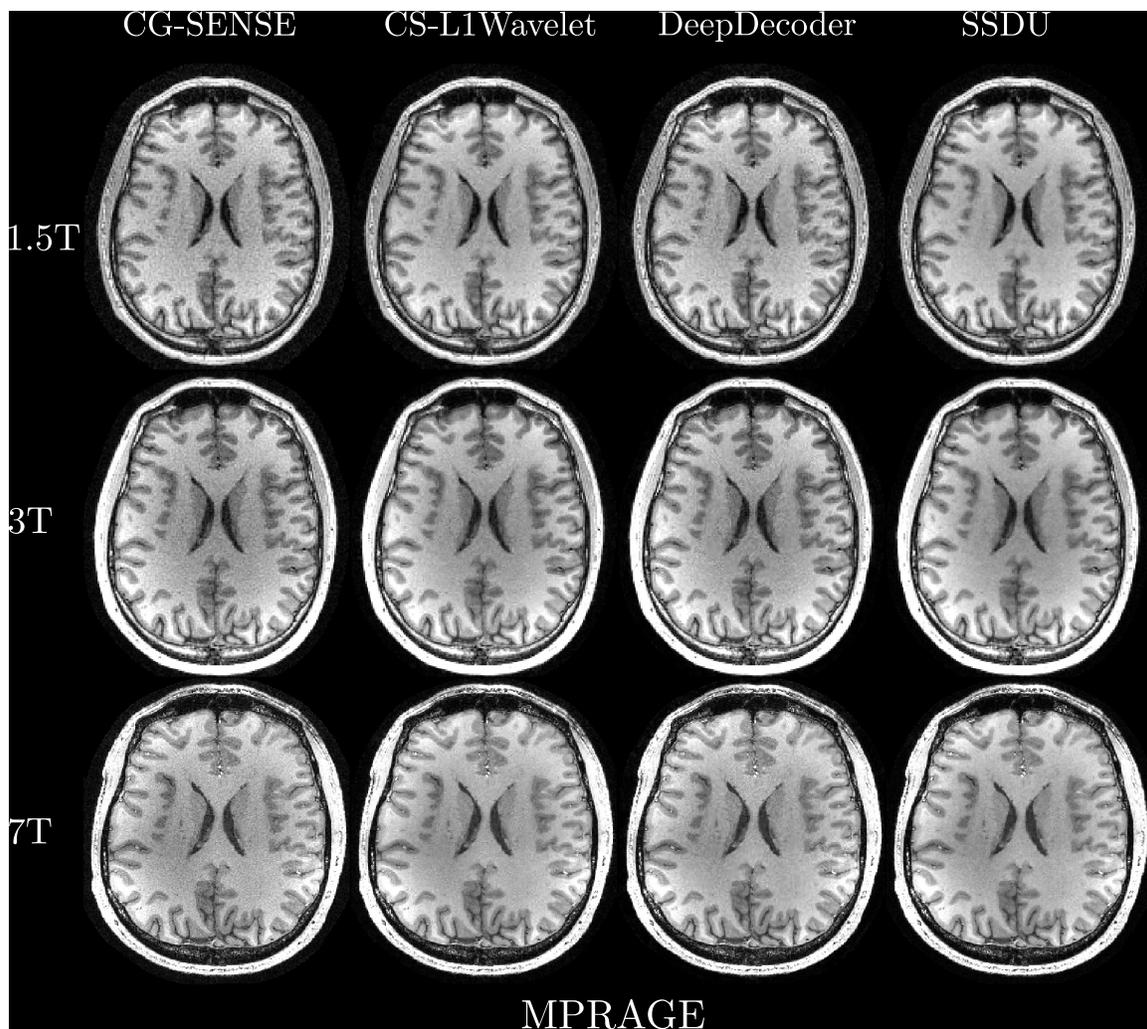


Figure 4: Axial slices from prospective reconstructions of MPRAGE scans of the brain at different field strengths. **Images are not co-registered**; The interpolation of image co-registration introduces blurring and thus was omitted. We chose slices at similar locations for visualization. CG-SENSE produces noisy but sharp reconstructions, and DeepDecoder produces smoother reconstructions with spatially varying noise and oversmoothing. CS-L1Wavelet and SSDU produce similarly smooth/sharp reconstructions.

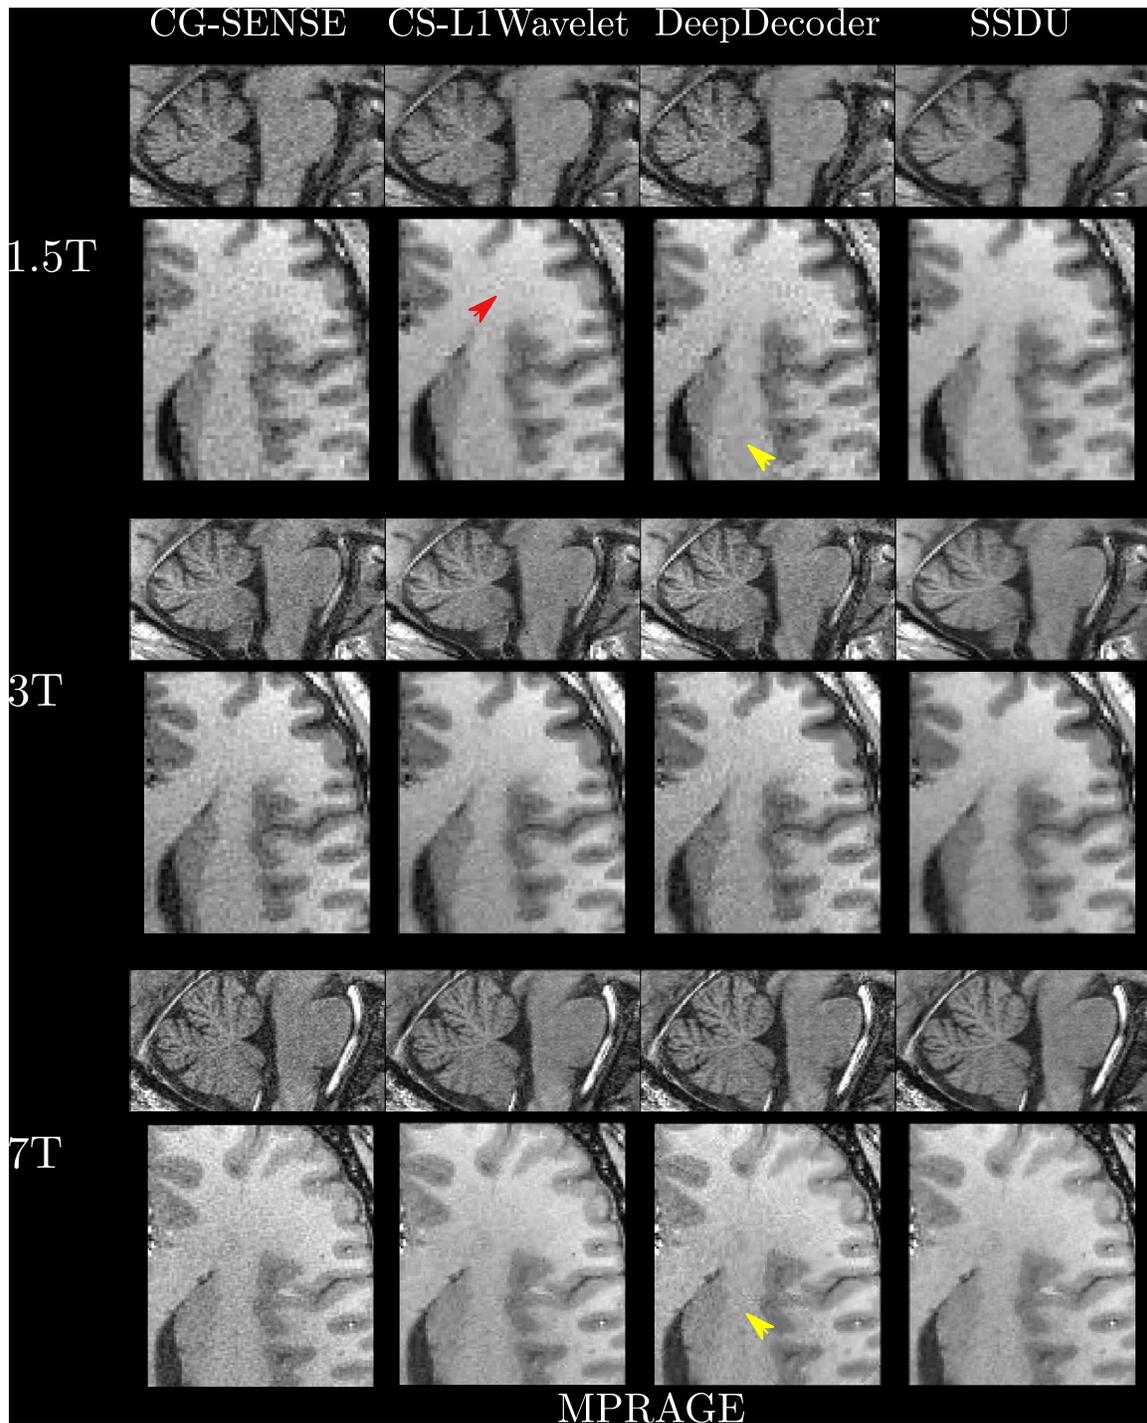


Figure 5: Closeups of the prospective reconstructions of MPRAGE scans of the brain at different field strengths; we show closeups of the cerebellum in a sagittal view as well as the left frontal lobe in an axial view. In the axial closeups, the spatially varying smoothness of DeepDecoder is apparent (yellow arrows); furthermore, wavelet artifacts of CS-L1Wavelet can be seen in, for example, the axial closeup at 1.5T (red arrow). In general, we can see that all methods improve in sharpness (as can be seen from the closeups of the corpus callosum) with increasing field strength.

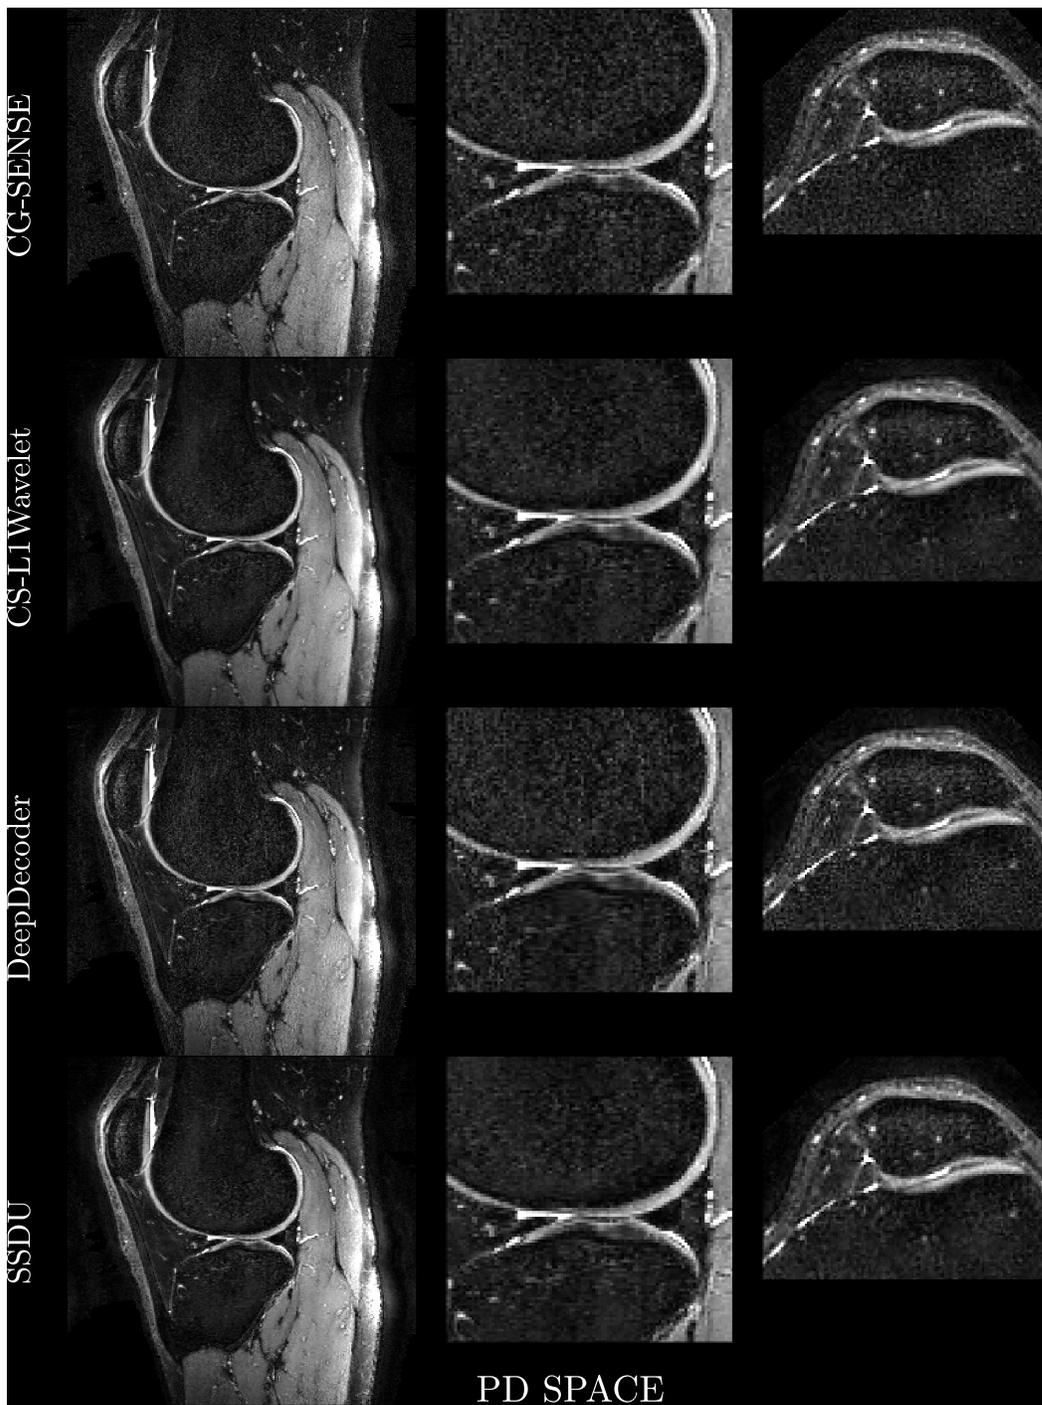


Figure 6: Prospective reconstructions from PD SPACE scans of the knee, where we show a sagittal slice as well as closeups on the articular cartilage interface in sagittal (femur) and axial (patella) views. Qualitatively, the main differences are between CS-L1Wavelet/SSDU and CG-SENSE/DeepDecoder, where the former group removes noise better than the latter.



Figure 7: Barplot of the no-reference image metrics averaged over all the subjects/different sequences in the generalizability study (top row). The arrow next to each metric indicates whether higher/lower scores are better. Barplots of the qualitative rating done by the MR scientists (pooled together) and the radiologist respectively (bottom row). We found statistically significant differences between all methods with respect to the no-reference image metrics. With respect to the MR scientists the following differences **were not** statistically significant: (DeepDecoder,CS-L1Wavelet,Sharpness),(DeepDecoder,SSDU,Sharpness), all of the aliasing comparisons, and (CS-L1Wavelet, SSDU, Overall Quality). All other comparisons were found to be statistically significant. With respect to the radiologist, all sharpness and aliasing comparisons were found to **not be** statistically significant. In the SNR comparisons, only (CG-SENSE/DeepDecoder vs. SSDU) were found to **be** statistically significant. In the overall quality comparisons, only (CG-SENSE vs DeepDecoder) and (CS-L1Wavelet vs. SSDU) were found to **not be** statistically significant. Overall, the no reference image metrics and human rating agree that CS-L1Wavelet/SSDU exhibit better overall image quality than DeepDecoder/CG-SENSE.

5. Discussion

In contrast to the previous literature, this work critically examines the validation and generalizability of self-supervised algorithms for undersampled MRI reconstruction through novel experiments with a focus on prospective reconstructions, the clinically relevant scenario. To this end, we analyze results from acquiring both fully-sampled and prospectively accelerated data on two phantoms and prospectively accelerated, in-vivo data over a wide variety of different sequences.

5.1 Validation using Prospectively Accelerated and Fully Sampled Data

Concerns about the differences between prospective and retrospective reconstructions were also raised in (Muckley et al., 2021), in the context of end-to-end, supervised methods for parallel MR image reconstruction. In particular, they noted that retrospective undersampling neglects potential differences in signal relaxation across echo trains, and verification should be performed before clinical use. From our results using both fully sampled and prospectively accelerated data, it is clear that for the 3D MPRAGE sequence, prospective vs. retrospective reconstructions can differ meaningfully, with retrospective reconstructions having greater fidelity to the fully sampled reconstruction; prospective reconstructions exhibit spatial distortions and local changes in contrast, with respect to the ground truth. This is despite the methods being tuned/trained on prospectively accelerated data; hence, this can be attributed to the differences in the prospectively vs. retrospectively sampled k-space data, potentially due to the different gradient patterns used in the sequences. This difference is relevant both for self-supervised and supervised machine learning methods; indeed, end-to-end, supervised methods which are trained on retrospective data may yield even greater distortion than self-supervised methods when prospective data is used for inference. However, the performance ranking of the different methods was the same in both prospective and retrospective reconstructions. Therefore, retrospective image quality cannot necessarily be taken as a reliable proxy for prospective image quality; however, they can be used to show differences between methods.

The quantitative results in the phantom show how ranking by PSNR and SSIM can be misleading, as images that are perceptually/qualitatively more similar to the ground truth (SSDU,CS-L1Wavelet) can have significantly worse or almost identical mean PSNR/SSIM scores than images which are less qualitatively similar (CG-SENSE,DeepDecoder). In contrast, ranking with the PercDis score, which measures distances between the feature activations within a pretrained classification network of the images rather than the images themselves, better matches with the perceptual quality of the images, showing that SSDU or SSDU/CS-L1Wavelet are better, by a significant margin (relatively with respect to the same differences in PSNR/SSIM), than the other methods. The PercDis score or perceptual loss (Johnson et al., 2016) was created precisely because they found this metric better suited for measuring perceptual similarity than PSNR/SSIM. This apparent trade-off between PSNR/SSIM and perceptual similarity is well-known in the computer vision community, where it is called the perception-distortion tradeoff (Blau and Michaeli, 2018). This concept has also recently been explored in MR; in (Adamson et al., 2021), the authors train an in-painting network on the Fastmri dataset, and use the features of intermediate layers for quantitative evaluation, producing a perceptual distance tailored for MR images.

In (Wang et al., 2019), the authors propose a new reconstruction method which uses distances in feature space (trained from ground truth MR reconstructions) to better recover textures/perceptual appearance than using just pixel-wise metrics.

5.2 Generalizability

We note that as our generalizability study is conducted on prospective reconstructions, which we showed can exhibit distortions relative to fully-sampled reconstructions, it cannot be considered as clinical validation; however, as all methods are affected the same way, this study still can give a good idea of how well each method generalizes. While one might conjecture that generalizability is less of a problem for self-supervised methods, if the parameters/hyperparameters of the methods are tuned for a specific sequence/anatomy as in our case, this could potentially impact the robustness of the methods, as these parameters/hyperparameters are obtained from training/tuning on 3D, brain MPRAGE scans acquired at 3T. This is despite the data consistency inherently embedded in CS-L1Wavelet, DeepDecoder, and SSDU.

Generalizability and robustness of reconstruction methods have been studied in the context of end-to-end, supervised methods for MR reconstruction in (Knoll et al., 2019; Hammernik et al., 2021; Antun et al., 2020). We briefly summarize some relevant conclusions from these articles. (Knoll et al., 2019) found that that different domain shifts reduced performance more than others (e.g. changing SNR vs. image contrast), and that transfer learning is a viable strategy for handling distribution shifts. (Hammernik et al., 2021) found that data consistency is important for robustness, and that at acceleration factor 4, distribution shifts are less of an issue. (Antun et al., 2020) found that supervised methods are vulnerable to adversarial perturbations, i.e. perturbations constructed such that minimal changes in the input data result in significant changes in the output. In (Darestani et al., 2021), the authors examine the robustness of end-to-end methods, compressed sensing, and variations of Deep Image Prior/DeepDecoder to distribution shifts, adversarial perturbations, and recovery of small features. They found that for both supervised and self-supervised methods, distribution shifts resulted in decreased PSNR/SSIM scores; in addition, the decrease was roughly the same for each method, preserving the ranking of the methods. Finally they found that all methods, including self-supervised methods, were vulnerable to adversarial attacks, including CS-L1Wavelet and DeepDecoder. Furthermore, Zhang et al. showed the vulnerability of SSDU to adversarial attacks, showing that this was primarily due to the data consistency term. Thus, CG-SENSE can also be assumed to be vulnerable. Therefore, all the methods used in this paper have been shown to be vulnerable to adversarial attacks. We note that these works are based on retrospective reconstructions/retrospective sampling from fully-sampled datasets for their validation.

In line with (Knoll et al., 2019), we found that different distribution shifts affected generalization differently; changing anatomy/contrast worsened the overall image quality rating in comparison to changing the field strength for all methods according to the MR scientists; in contrast, the radiologist found that only SSDU/CS-L1Wavelet worsened. However, as the mean scores in the knee scans for CG-SENSE/DeepDecoder were already 4 (the worst score), the decrease may not reflect any substantial difference in quality. As in (Hammernik et al., 2021), data consistency is crucial for the robustness of self-supervised methods as network

parameters are trained solely through the modelling/the acquired undersampled data; in particular, we do not see any hallucination that can occur with end-to-end networks without data consistency. Furthermore, we see that as CG-SENSE produces a plausible image with acceleration factor 5, this can explain why distribution shifts were not so troublesome, as the self-supervised methods mainly needed to denoise, rather than recover anatomy/missing high frequency details.

In contrast to (Darestani et al., 2021), our PSNR/SSIM results on the phantoms do not preserve the ranking between methods, although the PercDis results do, approximately. However, the qualitative metrics between distribution shifts over the different brain/knee scans seem to preserve ranking according to the no-reference image metrics/human ratings; this is consistent with PercDis being a better measure for perceptual image quality/similarity than PSNR/SSIM. In addition, the distribution shift in (Darestani et al., 2021) was between two, similar datasets of knee MRI, as compared to our distribution shifts, where we change anatomy, contrast, etc.

For a clinical scenario, it was of interest to see if self-supervised methods could potentially work, without retraining, on other sequences, as retraining after deployment could be impractical. Furthermore, while adversarial perturbations are valuable for studying the input stability of reconstruction methods, they need to be manually constructed for each method and added to the input data. As clinical MR reconstruction is a closed loop, this kind of manual perturbation would require hacking the internal MR computer. Therefore, transfer learning and adversarial perturbations were outside the scope of this work, although from (Hammernik et al., 2021; Knoll et al., 2019; Darestani et al., 2021), we would expect an increase in image quality from transfer learning and vulnerability to adversarial perturbations for the methods considered in this paper. For example, (Darestani and Heckel, 2021) found, in a retrospective study, that DeepDecoder had different optimal (judged by PSNR/SSIM) hyperparameters for brain vs. knee scans. However, SSDU and CS-L1Wavelet, tuned only on 3T MPRAGE brain data, are able to achieve an overall image quality of fair to good on a diverse dataset.

5.3 Ranking Methods and Quantitative Metrics

From a perceptual viewpoint (PercDis score, no-reference image metrics, human rating), SSDU and CS-L1Wavelet performed the best, with an edge to SSDU in the PercDis score/no-reference image metrics. From a pixel-wise metric viewpoint (PSNR,SSIM), DeepDecoder was better than or similar to all methods, as was also found in (Darestani et al., 2021). CG-SENSE consistently performed the worst or similarly to all methods over all metrics. With respect to validation, both approaches have their advantages and disadvantage; while pixel-wise metrics are the natural way to compare against a ground-truth, they may not correlate well with the perception of a radiologist. While perceptual metrics may be intuitive, the absence of ground truth can make it less objective. To our knowledge, current state of the art MR image reconstructions are generally not evaluated with perceptual metrics such as PercDis or (Adamson et al., 2021), which require ground truth, or the no-reference image quality metrics. However, given the close correspondence of the image quality metrics/PercDis to the human ratings/perceptual evaluation, as well as other evidence from

the literature (Woodard and Carley-Spencer, 2006; Zhang et al., 2018), perceptual metrics could be used as a complement to pixel-wise metrics/human ratings.

5.4 Implications for Future Methods and Validation

We note that while SSDU generally outperformed DeepDecoder, SSDU’s denoising network was trained on a dataset of 3T MPRAGE, thus learning a prior over multiple subjects. In contrast, DeepDecoder only learns/performs inference over a single slice at a time, thus limiting the amount of information in comparison to SSDU. In Korkmaz et al. (2022), the authors show that a Deep Image Prior based reconstruction can be fused explicitly with prior information from a dataset of fully sampled acquisitions to increase performance; such fusion of prior information could potentially also benefit DeepDecoder and other self-supervised methods which operate on a per slice basis. In addition, from the qualitative results, CS-L1Wavelet with regularization parameter tuned using the Noise2Self framework is competitive with SSDU. At lower acceleration factors, such as the one used in this paper, it is plausible that this result generalizes, such that compressed sensing reconstructions with optimally tuned regularization parameters can be competitive with state of the art machine learning methods, at least on a qualitative basis. For future validation, we conclude that appropriate regularization parameter tuning strategies should be used when comparing compressed sensing reconstructions to new methods. Finally, we note that as the theory behind the self-supervised methods we used (Deep Image Prior and blind denoising) form the basis for or are conceptually similar to many other self-supervised methods, it is plausible that the impressive robustness showed by these methods to a diverse range of realistic distribution shifts would generalize to future self-supervised methods.

5.5 Future of Validation

However, whatever metrics or datasets are used for validating methods, the ultimate test for reconstruction methods is the usefulness to radiologists for reliably diagnosing pathology in comparison to currently used methods (Recht et al., 2020; Roux et al., 2019). This can imply many things, including fine grained analysis of small textures/details/pathologies as well as tissue specific analysis, requiring novel datasets with extensive annotations by radiologists. (Zhao et al., 2021; Desai et al., 2021) are two recent works in this direction, providing datasets with bounding box annotations/pathology annotations to further validate reconstructions. To assist validating future methods, the datasets acquired for this paper will be made available online; see <https://www.melba-journal.org/papers/2022:022.html> for details.

6. Conclusion

Rigorous validation is required to introduce new reconstruction algorithms into clinical routines. In this study, validation of prospective reconstructions, generalizability, and different image quality metrics were investigated. The results show that self-supervised image reconstruction methods have potential, but that further development is required to not only improve image quality but also to define a reliable, standardized way of validating new

methods. Reliable validation can facilitate quicker translation to the clinical routine, with the ultimate goal of improving patient care.

Acknowledgments

This project is supported by the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie project TRABIT (agreement No 765148 to TY) and by the Swiss National Science Foundation (SNSF, Ambizione grant PZ00P2_185814 to EJC-R) We acknowledge access to the facilities and expertise of the CIBM Center for Biomedical Imaging, a Swiss research center of excellence founded and supported by Lausanne University Hospital (CHUV), University of Lausanne (UNIL), Ecole Polytechnique Fédérale de Lausanne (EPFL), University of Geneva (UNIGE) and Geneva University Hospitals (HUG).

Ethical Standards

The work follows appropriate ethical standards in conducting research and writing the manuscript, following all applicable laws and regulations regarding treatment of animals or human subjects.

Conflicts of Interest

Thomas Yu, Gian Franco Piredda, Gabriele Bonanno, Arun Joseph, Tom Hilbert and Tobias Kober are employed by Siemens Healthineers International AG, Switzerland.

References

- Philip M Adamson, Beliz Gunel, Jeffrey Dominic, Arjun D Desai, Daniel Spielman, Shreyas Vasanawala, John M Pauly, and Akshay Chaudhari. Ssf: Self-supervised feature distance as an mr image reconstruction quality metric. *NeurIPS 2021 Workshop on Deep Learning and Inverse Problems*, 2021.
- Mehmet Akçakaya, Burhaneddin Yaman, Hyungjin Chung, and Jong Chul Ye. Unsupervised deep learning methods for biological image reconstruction. *arXiv preprint arXiv:2105.08040*, 2021.
- Vegard Antun, Francesco Renna, Clarice Poon, Ben Adcock, and Anders C Hansen. On instabilities of deep learning in image reconstruction and the potential costs of ai. *Proceedings of the National Academy of Sciences*, 117(48):30088–30095, 2020.
- Joshua Batson and Loic Royer. Noise2self: Blind denoising by self-supervision. *Proceedings of the International Conference on Machine Learning*, pages 524–533, 2019.
- Yochai Blau and Tomer Michaeli. The perception-distortion tradeoff. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6228–6237, 2018.

- Robert Bridson. Fast poisson disk sampling in arbitrary dimensions. *SIGGRAPH sketches*, 10(1):1, 2007.
- Mohammad Zalbagi Darestani and Reinhard Heckel. Accelerated mri with un-trained neural networks. *IEEE Transactions on Computational Imaging*, 7:724–733, 2021.
- Mohammad Zalbagi Darestani, Akshay S. Chaudhari, and Reinhard Heckel. Measuring robustness in deep learning based compressive sensing. *International Conference on Machine Learning*, 139:2433–2444, 2021. URL <http://proceedings.mlr.press/v139/darestani21a.html>.
- Arjun D Desai, Andrew M Schmidt, Elka B Rubin, Christopher Michael Sandino, Marianne Susan Black, Valentina Mazzoli, Kathryn J Stevens, Robert Boutin, Christopher Re, Garry E Gold, et al. Skm-tea: A dataset for accelerated mri reconstruction with dense image labels for quantitative clinical evaluation. *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- K. Epperson, A.M Sawyer, M. Lustig, M.T. Alley, M. Uecker, P. Virtue, P. Lai, and Vasanaawala SS. Creation of fully sampled mr data repository for compressed sensing of the knee. *Proceedings of Society for MR Radiographers and Technologists, 22nd Annual Meeting. Salt Lake City, Utah, USA.*, 2013.
- Mark A Griswold, Peter M Jakob, Robin M Heidemann, Mathias Nittka, Vladimir Jellus, Jianmin Wang, Berthold Kiefer, and Axel Haase. Generalized autocalibrating partially parallel acquisitions (grappa). *Magnetic Resonance in Medicine*, 47(6):1202–1210, 2002.
- Kerstin Hammernik and Florian Knoll. Chapter 2 - machine learning for image reconstruction. In S. Kevin Zhou, Daniel Rueckert, and Gabor Fichtinger, editors, *Handbook of Medical Image Computing and Computer Assisted Intervention*, The Elsevier and MICCAI Society Book Series, pages 25–64. Academic Press, 2020. ISBN 978-0-12-816176-0. doi: <https://doi.org/10.1016/B978-0-12-816176-0.00007-7>. URL <https://www.sciencedirect.com/science/article/pii/B978012816176000077>.
- Kerstin Hammernik, Teresa Klatzer, Erich Kobler, Michael P Recht, Daniel K Sodickson, Thomas Pock, and Florian Knoll. Learning a variational network for reconstruction of accelerated mri data. *Magnetic Resonance in Medicine*, 79(6):3055–3071, 2018.
- Kerstin Hammernik, Jo Schlemper, Chen Qin, Jinming Duan, Ronald M. Summers, and Daniel Rueckert. Systematic evaluation of iterative deep neural networks for fast parallel mri reconstruction with sensitivity-weighted coil combination. *Magnetic Resonance in Medicine*, 86(4):1859–1872, 2021. doi: <https://doi.org/10.1002/mrm.28827>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.28827>.
- Reinhard Heckel and Paul Hand. Deep decoder: Concise image representations from untrained non-convolutional networks. *Proceedings of the International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=rylV-2C9KQ>.

- Reinhard Heckel and Mahdi Soltanolkotabi. Compressive sensing with un-trained neural networks: Gradient descent finds a smooth approximation. *Proceedings of the International Conference on Machine Learning*, pages 4149–4158, 2020.
- Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. *Proceedings of the European conference on computer vision*, pages 694–711, 2016.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Florian Knoll, Kerstin Hammernik, Erich Kobler, Thomas Pock, Michael P Recht, and Daniel K Sodickson. Assessment of the generalization of learned image reconstruction and the potential for transfer learning. *Magnetic Resonance in Medicine*, 81(1):116–128, 2019.
- Florian Knoll, Kerstin Hammernik, Chi Zhang, Steen Moeller, Thomas Pock, Daniel K Sodickson, and Mehmet Akcakaya. Deep-learning methods for parallel magnetic resonance imaging reconstruction: A survey of the current approaches, trends, and issues. *IEEE Signal Processing Magazine*, 37(1):128–140, 2020a.
- Florian Knoll, Jure Zbontar, Anuroop Sriram, Matthew J Muckley, Mary Bruno, Aaron Defazio, Marc Parente, Krzysztof J Geras, Joe Katsnelson, Hersh Chandarana, et al. fastmri: A publicly available raw k-space and dicom dataset of knee images for accelerated mr image reconstruction using machine learning. *Radiology: Artificial Intelligence*, 2(1): e190007, 2020b.
- Yilmaz Korkmaz, Salman UH Dar, Mahmut Yurt, Muzaffer Özbey, and Tolga Cukur. Unsupervised mri reconstruction via zero-shot learned adversarial transformers. *IEEE Transactions on Medical Imaging*, 2022.
- Jiaming Liu, Yu Sun, Cihat Eldeniz, Weijie Gan, Hongyu An, and Ulugbek S Kamilov. Rare: Image reconstruction using deep priors learned without groundtruth. *IEEE Journal of Selected Topics in Signal Processing*, 14(6):1088–1099, 2020.
- Andreas Lugmayr, Martin Danelljan, Radu Timofte, Namhyuk Ahn, Dongwoon Bai, Jie Cai, Yun Cao, Junyang Chen, Kaihua Cheng, Se Young Chun, Wei Deng, Mostafa El-Khamy, Chiu Man Ho, Xiaozhong Ji, Amin Kheradmand, Gwantae Kim, Hanseok Ko, Kanghyu Lee, Jungwon Lee, Hao Li, Ziluan Liu, Zhi-Song Liu, Shuai Liu, Yunhua Lu, Zibo Meng, Pablo Navarrete Micheli, Christian Micheloni, Kalpesh Prajapati, Haoyu Ren, Yonghyeok Seo, Wan-Chi Siu, Kyung-Ah Sohn, Ying Tai, Rao Muhammad Umer, Shuangquan Wang, Huibing Wang, Timothy Haoning Wu, Haoning Wu, Biao Yang, Fuzhi Yang, Jaejun Yoo, Tongtong Zhao, Yuanbo Zhou, Haijie Zhuo, Ziyao Zong, and Xueyi Zou. NTIRE 2020 challenge on real-world image super-resolution: Methods and results. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR Workshops*, pages 2058–2076, 2020. doi: 10.1109/CVPRW50498.2020.00255. URL https://openaccess.thecvf.com/content_CVPRW_2020/html/w31/Lugmayr_NTIRE_2020_Challenge_on_Real-World_Image_Super-Resolution_Methods_and_Results_CVPRW_2020_paper.html.

- Michael Lustig, David Donoho, and John M Pauly. Sparse mri: The application of compressed sensing for rapid mr imaging. *Magnetic Resonance in Medicine*, 58(6):1182–1195, 2007.
- Anish Mittal, Anush K Moorthy, and Alan C Bovik. Blind/referenceless image spatial quality evaluator. *Proceedings of the Forty Fifth ASILOMAR Conference on Signals, Systems and Computers*, pages 723–727, 2011.
- Matthew J Muckley, Bruno Riemenschneider, Alireza Radmanesh, Sunwoo Kim, Geunu Jeong, Jingyu Ko, Yohan Jun, Hyungseob Shin, Dosik Hwang, Mahmoud Mostapha, et al. Results of the 2020 fastmri challenge for machine learning mr image reconstruction. *IEEE Transactions on Medical Imaging*, 40(9):2306–2317, 2021.
- Emilie Mussard, Tom Hilbert, Christoph Forman, Reto Meuli, Jean-Philippe Thiran, and Tobias Kober. Accelerated mp2rage imaging using cartesian phyllotaxis readout and compressed sensing reconstruction. *Magnetic Resonance in Medicine*, 84(4):1881–1894, 2020.
- Frank Ong and Michael Lustig. Sigpy: a python package for high performance iterative reconstruction. *Proceedings of the International Society of Magnetic Resonance in Medicine, Montréal, QC*, 4819, 2019.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32:8026–8037, 2019.
- Klaas P Pruessmann, Markus Weiger, Markus B Scheidegger, and Peter Boesiger. Sense: sensitivity encoding for fast mri. *Magnetic Resonance in Medicine*, 42(5):952–962, 1999.
- Klaas P Pruessmann, Markus Weiger, Peter Börnert, and Peter Boesiger. Advances in sensitivity encoding with arbitrary k-space trajectories. *Magnetic Resonance in Medicine*, 46(4):638–651, 2001.
- Michael P Recht, Jure Zbontar, Daniel K Sodickson, Florian Knoll, Nafissa Yakubova, Anuroop Sriram, Tullie Murrell, Aaron Defazio, Michael Rabbat, Leon Rybak, et al. Using deep learning to accelerate knee mri at 3 t: results of an interchangeability study. *American journal of Roentgenology*, 215(6):1421, 2020.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- Marion Roux, Tom Hilbert, Mahmoud Hussami, Fabio Becce, Tobias Kober, and Patrick Omoumi. Mri t2 mapping of the knee providing synthetic morphologic images: comparison to conventional turbo spin-echo mri. *Radiology*, 293(3):620–630, 2019.
- David Salomon. *Data compression: the complete reference*. Springer Science & Business Media, 2004.

- Jian Sun, Huibin Li, Zongben Xu, et al. Deep admn-net for compressive sensing mri. *Advances in neural information processing systems*, 29, 2016.
- Martin Uecker, Peng Lai, Mark J Murphy, Patrick Virtue, Michael Elad, John M Pauly, Shreyas S Vasanawala, and Michael Lustig. Espirit—an eigenvalue approach to autocalibrating parallel mri: where sense meets grappa. *Magnetic Resonance in Medicine*, 71(3): 990–1001, 2014.
- Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9446–9454, 2018.
- N Venkatanath, D Praneeth, Maruthi Chandrasekhar Bh, Sumohana S Channappayya, and Swarup S Medasani. Blind image quality evaluation using perception based features. *Proceedings of the Twenty First National Conference on Communications (NCC)*, pages 1–6, 2015.
- Ke Wang, Jonathan I. Tamir, and Stella X. Yu. High-fidelity reconstruction with instance-wise discriminative feature matching loss. *Proc. Intl. Soc. Mag. Reson. Med.* 28, 2019.
- Zhou Wang, Hamid R Sheikh, and Alan C Bovik. No-reference perceptual quality assessment of jpeg compressed images. *Proceedings of the International Conference on Image Processing*, 1:I–I, 2002.
- Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
- Jeffrey P Woodard and Monica P Carley-Spencer. No-reference image quality metrics for structural mri. *Neuroinformatics*, 4(3):243–262, 2006.
- Burhaneddin Yaman, Seyed Amir Hossein Hosseini, Steen Moeller, Jutta Ellermann, Kâmil Uğurbil, and Mehmet Akçakaya. Self-supervised learning of physics-guided reconstruction neural networks without fully sampled reference data. *Magnetic Resonance in Medicine*, 84(6):3172–3191, 2020.
- Chi Zhang, Jinghan Jia, Burhaneddin Yaman, Steen Moeller, Sijia Liu, Mingyi Hong, and Mehmet Akçakaya. Instabilities in conventional multi-coil mri reconstruction with small adversarial perturbations. In *2021 55th Asilomar Conference on Signals, Systems, and Computers*, pages 895–899. IEEE.
- Zhicheng Zhang, Guangzhe Dai, Xiaokun Liang, Shaode Yu, Leida Li, and Yaoqin Xie. Can signal-to-noise ratio perform as a baseline indicator for medical image quality assessment. *IEEE Access*, 6:11534–11543, 2018.
- Ruiyang Zhao, Burhaneddin Yaman, Yuxin Zhang, Russell Stewart, Austin Dixon, Florian Knoll, Zhengnan Huang, Yvonne W Lui, Michael S Hansen, and Matthew P Lungren. fastmri+: Clinical pathology annotations for knee and brain fully sampled multi-coil mri data. *arXiv preprint arXiv:2109.03812*, 2021.

Appendix

Hyperparameter Tuning

For example, to set the regularization parameter of CS-L1Wavelet, we treat it as a function with a single parameter (λ). We can then optimize this parameter using the Noise2Self training framework to estimate the λ which minimizes the noise-free error between simulated measurements and the acquired measurements. Concretely, we fix 20 logarithmically spaced values from 0.00001 to 0.1. We set each value as λ and run 50 image reconstructions corresponding to different, random masks and average the corresponding errors with respect to the complementary mask in order to approximate the true measurement error associated with using each value. We then select the value with the lowest measurement error as the optimal regularization parameter. This is done for each slice in each subject; the final regularization value which is used throughout this paper is the average over all subjects. The hyperparameters of DeepDecoder and SSDU are set similarly with a grid search over the network hyperparameters, albeit over a much smaller set of data due to the high computational demand.

	1.	2.	3.	4.	5.	6.	7.	8.	9.	10.
Sequence Type	MPRAGE	MPRAGE	MPRAGE	MPRAGE	MPRAGE	MPRAGE	SPACE	SPACE	SPACE	SPACE
Field Strength (T)	1.5	3	7	3	3	3	3	3	3	3
Body Part	Brain	Knee	Knee							
Coils	1Tx/20Rx	1Tx/64Rx	8pTx/32Rx	1Tx/20Rx	1Tx/64Rx	1Tx/64Rx	1Tx/64Rx	1Tx/64Rx	1Tx/18Rx	1Tx/18Rx
Resolution (mm³)	1.3x1.3x1.2	1x1x1	0.7x0.7x0.7	1x1x1	1x1x1	1x1x1	1x1x1	1x1x1	0.3x0.3x0.6	0.3x0.3x0.6
Field of View (mm³)	240x240x160	256x240x208	250x219x179	256x240x208	256x240x208	256x240x208	250x250x176	250x250x176	160x160x134	160x160x115
Inversion Time (s)	1	0.9	1.1	0.9	0.9	0.972	-	2.05	-	-
Repetition Time (s)	2.4	2.3	2.5	2.3	2.3	1.93	0.7	7	0.9	1
Echo Time (ms)	3.47	2.9	2.87	2.9	2.9	2.61	11	392	29	108
Echo Spacing (ms)	7.86	6.88	7.8	6.88	6.88	6.28	3.72	3.66	4.84	5.12
Bandwidth (Hz/Px)	180	240	250	240	240	280	630	651	488	416
Turbo Factor	192	198	250	198	198	198	42	220	35	44
Acceleration Factor	4.2	5	5	5	5	5	4	6	7	7
Acquisition Time	1:28 min	1:34 min	2:42 min	1:34 min	1:34 min	1:20 min	3:27 min	3:46 min	4:41 min	3:52 min

Table 2: Detailed sequence parameters of all used datasets. We note that spiral phyllotaxis sampling Mussard et al. (2020) and Poisson disk sampling Bridson (2007) were used for the MPRAGE and SPACE sequences respectively.

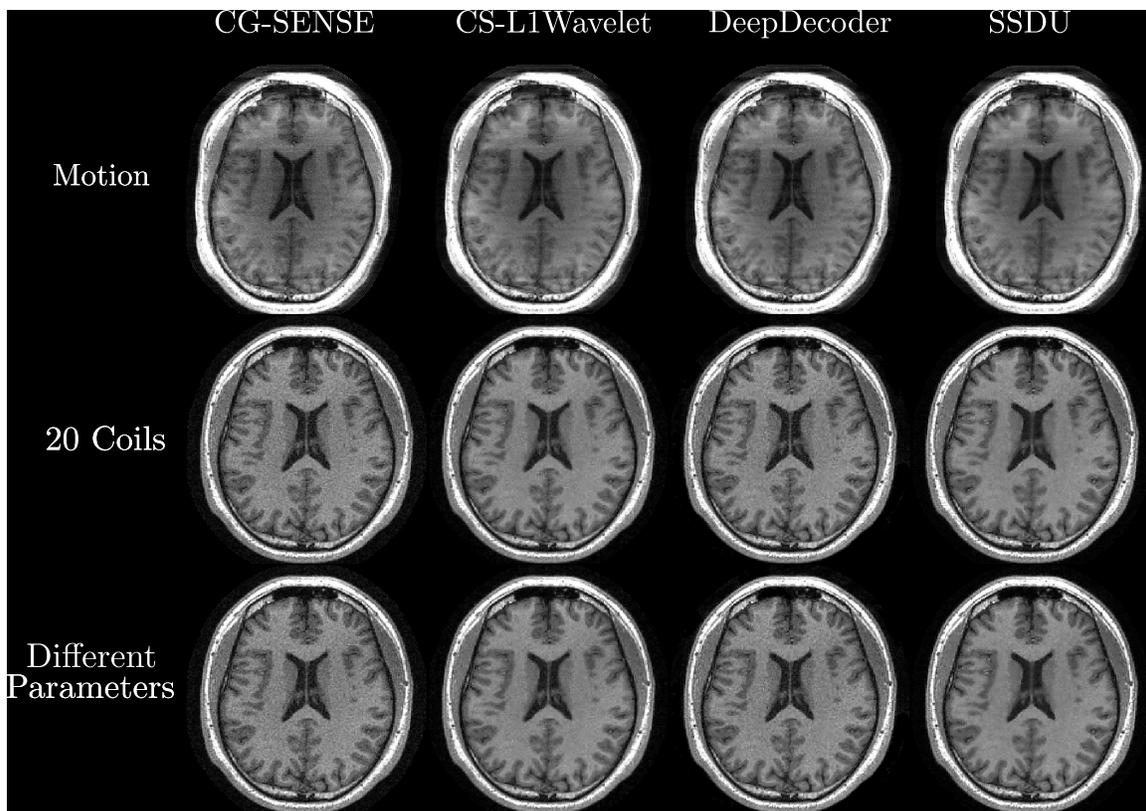


Figure 8: Here we show axial brain slice reconstructions from three different perturbations of the MPRAGE sequence: the addition of motion, using 20 coils instead 64 coils, and changing the parameters of the MPRAGE sequence. The images are **not** registered due to interpolation effects from co-registration.

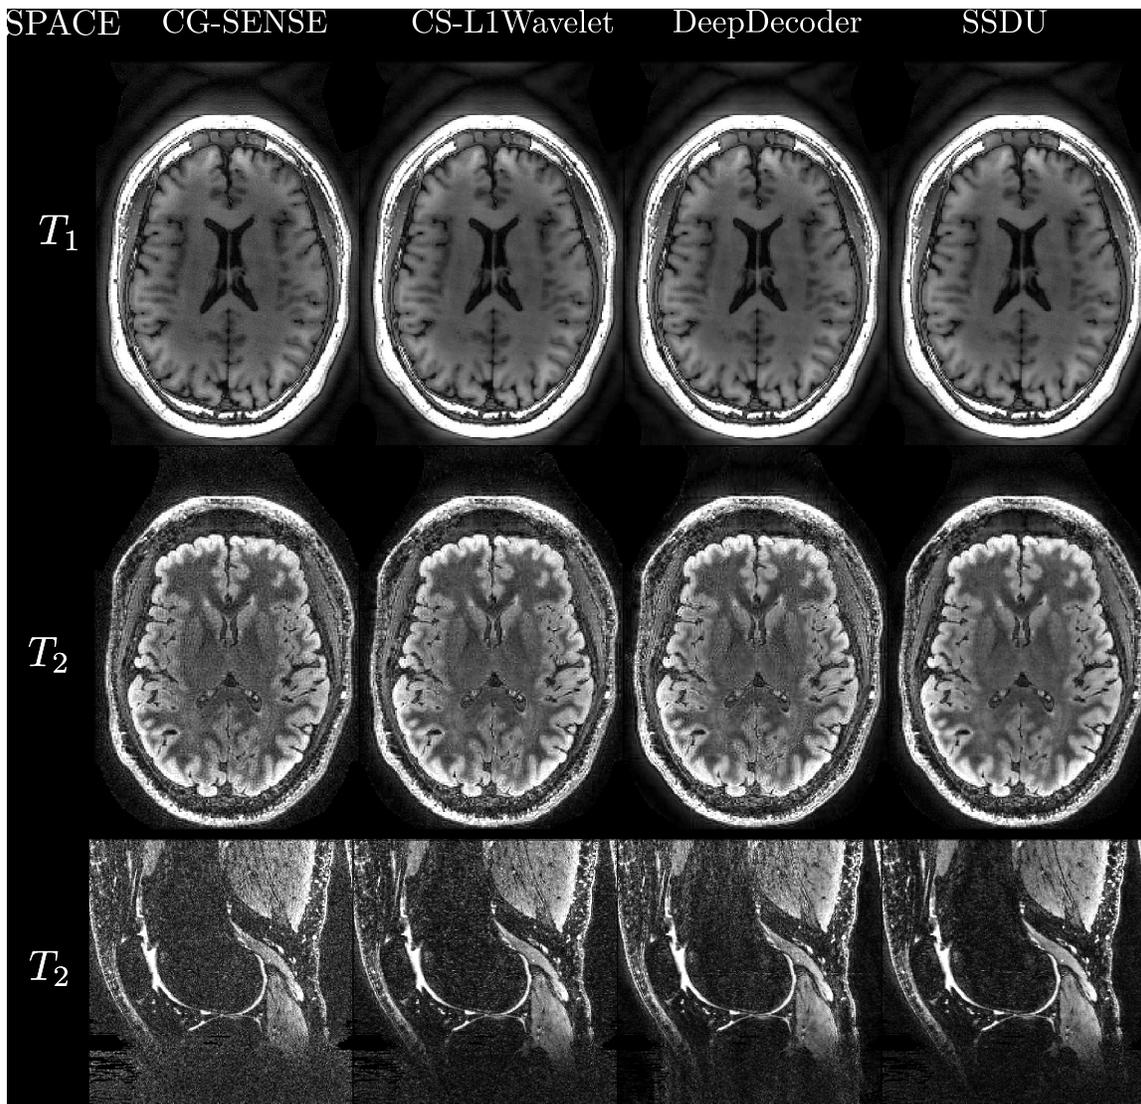


Figure 9: Here we show axial brain slices and a sagittal knee slice from the reconstructions from the SPACE acquisitions. The images are **not** registered due to interpolation effects from co-registration.