SANGRIA: one-Shot leArNinG super-ResolutIon with Adversarial training for accelerated Magnetic Resonance Imaging

Swetali Nimje, Ludovic de Rochefort, Thierry Artières

▶ **To cite this version:**

Swetali Nimje, Ludovic de Rochefort, Thierry Artières. SANGRIA: one-Shot leArNinG super-ResolutIon with Adversarial training for accelerated Magnetic Resonance Imaging. JOINT CAp (Conférence sur l'Apprentissage automatique) and RFIAP (Reconnaissance des Formes, Image, Apprentissage et Perception, Jul 2022, VANNES, France. hal-04050162

HAL Id: hal-04050162
https://hal.science/hal-04050162

Submitted on 29 Mar 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

SANGRIA: one-Shot leArNinG super-ResolutIon with Adversarial training for accelerated Magnetic Resonance Imaging

Nimje S[*1,2], de Rochefort L[2], and Artières T[†1]

[1]Laboratoire d'Informatique et Systèmes UMR 7020 Aix-Marseille Univ., CNRS, Ecole Centrale de Marseille, Faculté des Sciences, 163 avenue de Luminy, Marseille, France.
[2]Aix-Marseille Univ., CNRS, CRMBM, UMR7339, Faculté de Médecine, 27 boulevard Jean Moulin, Marseille, France.

April 13, 2022

Abstract

Magnetic Resonance Imaging (MRI) acquisition is performed sequentially in the spatial-frequency domain (k-space) and involves several views of the object/subject. To accelerate the acquisition, k-space lines are often undersampled on a cartesian grid. Parallel imaging reconstruction algorithms are then applied to recover unseen lines. We consider this super-resolution problem in k-space to further reduce the acquisition time with deep learning and to reduce the costs associated with this expensive medical imaging technology. Because the sensors are specific to anatomical regions and experimental setups, it pushes toward learning a reconstruction model on a per-image basis, i.e. for every scan and image. Here, we propose an extension of state-of-the-art MRI reconstruction methods where the super-resolution task in k-space is solved with a convolutional neural network, and where an adversarial strategy using a patch discriminator in image space is used to reach higher undersampling rates. Both parts are trained in a one-shot learning setting. It is demonstrated both using simulated and in vivo brain experiments that this combined approach provides enhanced image quality for undersampling rates larger than the ones used in a clinical routine on multi-slice 2D T2-weighted imaging sequences, making the approach readily applicable as an alternative reconstruction strategy in MRI systems for 2D parallel imaging, and which could be further extended to accelerate 3D imaging sequences.

*email: swetali.NIMJE@univ-amu.fr
†Corresponding author, thierry.artieres@lis-lab.fr

Keywords: Super-resolution, convolutional neural networks, adversarial, one-shot learning, MRI, parallel imaging, k-space, undersampling, neuroimaging.

1 Introduction

Magnetic resonance imaging (MRI) is a major diagnostic tool indicated in many diseases and is central to biomedical research. To provide a few examples, it is used in oncology to characterize tumour evolution, response to treatments or surgical planning; in neurology to detect and follow up brain lesions such as in multiple sclerosis or epilepsy, to locate fine structures deep in the brain for treating Parkinson's diseases or to characterize atrophy in Alzheimer's disease. MRI is extremely versatile in the information it provides thanks to a multiplicity of contrast mechanisms. However, this is an expensive technology, and MRI exams are long, typically lasting on the order of 30-45 minutes. Reducing the duration of an MRI exam while keeping the same information is challenging, but it is a key societal issue that would enable a wider availability and/or higher diagnostic value.

There is a series of operations that are performed from the raw MR signals to the final images that are used for diagnosis. Indeed, MRI scanning principles are based on the sequential acquisition of raw k-space data which corresponds to the spatial frequencies of the object/subject under observation. These sampled signals are complex-valued by nature. Additionally, the object is simultaneously seen by several sensors, called receiver coils. From a fully-sampled k-space, an MR image is reconstructed by applying an inverse Fourier

transform, and images are then combined from the different coils with, e.g., a sum-of-squares combination [1].

The most common way to reduce the acquisition time is to acquire fewer measurements by undersampling k-space, and exploit the redundancy coming from the multiple receiver-coil views to recover the missing k-space lines. This strategy is called parallel imaging (PI), which has been implemented in clinical MRI systems since its development in the years 2000's [2], [3]. It corresponds to a super-resolution problem in k-space. A standard reconstruction algorithm is GRAPPA [3].

A recent class of MR reconstruction techniques are based on deep learning (DL) [4]–[7]. AUTOMAP [5] performs the reconstruction tasks based on extensive simulations of the acquisition processes, while other approaches are more directly focusing on the super-resolution problem [4]. In order to perform this task, variational networks [8] or generative models [9], [10] were also proposed. DL based approaches hold great promises in further improving the reconstruction quality, however, it often requires a large amount of data for training [11]. Additionally, DL models trained on one type of data might not be adapted to the variety of MRI experimental conditions. Indeed, DL based approaches pose a challenge in reconstructing high-quality images for PI because of different acquisition parameters, anatomical territories and configuration of the receiver coils, which impede or reduce the capacity of DL reconstruction approaches to be generally applicable to all MRI acquisitions. Additionally, models trained using large datasets may also fail to generalize on rare pathologies, and this is one major concern in the trust in DL-based diagnostic tools and for their certification [12].

For the specific task of super-resolution in k-space, a family of approaches have focused on learning simple models with few parameters enabling learning a new model dedicated to each single image [6]. The goal of recovering a High-Resolution (HR) fully sampled k-space from the low resolution under-sampled k-space data can be formalized as a Single Image Super-Resolution (SISR) for MRI. In computer vision literature, SISR within the GAN framework [13] has gain popularity (e.g.[14]). This idea has the advantage to require less data for training, including scan-specific ones, but does not make use of adversarial learning to enhance the reconstructed image quality. This approach was proposed in computer vision using patch-discriminator [15] and further developed for image translation tasks. A number of studies have since been proposed to exploit adversarial training for MRI reconstruction. Wang et al. [14] proposed a GAN framework using a patch-discriminator for 3D single image super-resolution on T1-weighted brain MRI images.

In this work, we focus on uniform undersampled acquisitions for which we propose a novel one-shot super-resolution method that combines parallel imaging-inspired DL techniques following the line of works opened by [3], [6]. A convolutional neural network operating in the k-space domain is combined with patch-based adversarial learning in the image domain and trained for every new single image. It is demonstrated both in simulation and in vivo that this approach can achieve higher visual quality at higher acceleration rates than current standard MRI reconstructions of undersampled 2D acquisitions of brain images. The adversarial discriminator does bring improvements in reconstructed images quality although this does not translate into objective measurements and on the generalization ability of learned models. Our approaches could be extended in order to further accelerate the acquisition of 3D MRI k-space data.

2 Background and Related works

2.1 Background

In MRI, the signals collected are inherently complex-valued. The MRI signal measured by the i^{th} coil is acquired in k-space and given by

$$X_i = \mathcal{F}(c_i x), i \in 1, 2, ..n_c \tag{1}$$

where $\mathcal{F}$ is the Fourier transform, c_i is a complex-valued diagonal matrix encoding sensitivity of the i^{th} coil, n_c is the number of coils, and x is the image to estimate. In parallel imaging (PI), k-space is undersampled by

$$S_i = M_u X_i + \mathcal{N}(0, \sigma^2), i \in 1, 2, ..n_c \tag{2}$$

where M_u is a binary mask operator and $\mathcal{N}(0, \sigma^2)$ is the measurement noise that follows a normal distribution on both real and imaginary part with zero mean and a variance σ^2, assumed to follow the same statistics for all coils. The super-resolution problem can be formalized as estimating a fully-sampled version of the signals from the under-sampled multicoil signals. This fully-sampled version can be writen as:

$$\hat{S}_i = M_f X_i, i \in 1, 2, ..n_c \tag{3}$$

where M_f is a binary mask that samples k-space sufficiently such that an inverse Fourier transform provides image estimates free of aliasing. When a fully-sampled k-space is acquired or predicted, a single combined image is reconstructed using, e.g, the sum-of-squares (SoS) over coils as

$$\hat{x}_{SoS} = \sqrt{\sum_{i=1}^{n_c} \mathcal{F}^{-1}(\hat{S}_i)^2} \tag{4}$$

where $\mathcal{F}^{-1}$ is the inverse Fourier transform. In the following section, the relevant convolutional models to perform the upsampling task are detailed.

2.2 Related works

Generalized Auto-calibrating Partially Parallel Acquisitions (GRAPPA) have been broadly used in clinical Magnetic Resonance Imaging. GRAPPA estimates the missing k-space lines in each coil by a linear combination of its neighbourhood acquired data overall coils. During acquisition, the central region of k-space (which corresponds to low spatial frequencies) is fully sampled. This region, called the Auto-Calibration Signal or ACS, is used to estimate the GRAPPA weights G. The under-sampling is simulated in the ACS in a uniform manner as shown in the figure 1. The GRAPPA weights are then used at all spatial locations, which can be formalised as a complex-valued convolution in k-space from n_c channels to n_c channels.

The recently-proposed RAKI [6] non-linearly estimates the missing k-space lines using multiple convolutional layers and non-linear activation layers Rectified Linear Unit (ReLU). The reconstruction is similar to GRAPPA but uses a deep convolutional network instead of linear convolutional kernels for interpolation in k-space. The initial RAKI implementation suffered from longer run time because of large number of convolution networks were required. An alternate approach called Line-By-Line (LBL) [16] CNN for RAKI was introduced for clinically acceptable computational time while preserving the reconstructed image quality. Here, CNN models in this line are further developed and trained in an adversarial way.

3 Problem setting

In MRI, k-space data can be acquired in 2D or 3D. Here, we consider 2D acquisitions of several contiguous slices, which is called 2D multi-slice imaging. In the acquisition process, after a slice-selection, signals only originate from the selected slice and k-space data are acquired line by line in a sequential 2D scanning process. A line of the 2D k-space is fully-sampled along the frequency-encoding direction (the first dimension), but each time with different phase-encoding, i.e., with different sampling location in the second dimension. To get an idea, in all the experiments that we will consider, volumes are tensors of dimension n_s, n_c,n_x,n_y where n_s is the number of slices (2D images), n_c is the number of coils, n_x and n_y stand for the fully-sampled k-space sizes along frequency-encoding and phase-encoding directions, respectively.

According to the Nyquist–Shannon sampling theorem, k-space is fully sampled when the distance between two k-space samples is smaller than the inverse of the object size along this direction. In 2D acquisitions, it means that phase-encoding lines should not be too far appart to ensure that the reconstructed image via an inverse Fourier transform is free of aliasing . The acceleration rate, R, is defined as the ratio of the number of k-space lines that would be required to fullfill Nyquist-Shannon criterion, to the number of effectively acquired lines in an accelerated scan. The undersampling could either be uniform (every other R line in k-space), pseudo-random or ad-hoc. In this work, we focus on uniform undersampling. Slices are considered one at a time for learning a reconstrution model from an undersampled view of it based on the ACS, and the model is then applied on undersampled data to estimate fully-sampled k-spaces for this specific slice. In the next section is described how such a model is trained and regularized using an adversarial strategy.

4 Model

We propose an MRI reconstruction model driven by an adversarial loss as shown in Fig 2. We use as super-resolution model, G, state-of-the-art zero-shot reconstruction models such as GRAPPA and RAKI for uniform sampling. We perform Line-by-Line reconstruction as proposed in [16]. Such classical models did not include validation during the training stage which leads to overfitting. To overcome this issue, a train-validation partition scheme for early stopping is introduced as shown in Figure 1.

More specifically for a given acceleration rate R, we create $R-1$ networks which correspond to the $R-1$ lines to be predicted. From undersampling the ACS, we can obtain R data-target pairs that can be used as reference data. We thus use $R-1$ data-target pairs for training and the remaining pair for validation. To take

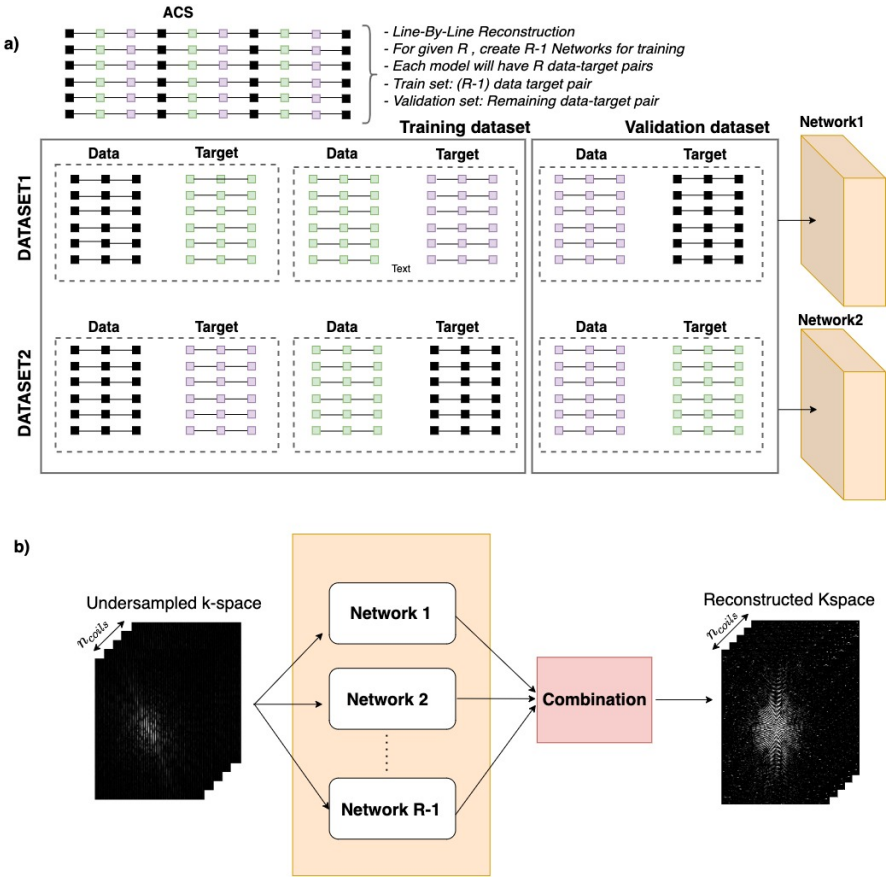

Figure 1: a) An overview of the proposed Train-Validation Partition for the super-resolution network for uniform sampling at R=3. Line-by-Line reconstruction is performed. For given R, R-1 networks are created for training. Each network learns the parameters from an undersampled view of the ACS and (b) After the training is performed, the weights are applied on the undersampled k-space data to obtain the final reconstruction i.e fully sampled kspace. it is a combination of the predicted lines from each network.

into account the complex nature of k-space data, we use complex implementation of the layers, as well as real implementations initialy reported in the literature [6] for comparison.

Moreover, to enhance the quality of the reconstructed images, an adversarial discriminator that operates on patches in image space (after inverse Fourier transform and coil combination) is proposed. The adversarial model is trained from a single fully-sampled coil-combined image. This patch-based discriminative network, D, is trained to distinguish between patches generated from the fully sampled image as well as from the reconstructed image. It is a convolutional neural network that takes a patch as input and outputs a classification score.

4.1 Model description

Super-resolution model. We use standard dimensioning of GRAPPA and RAKI models as used in previous studies [6]. Both models are convolutional networks. GRAPPA consists in one 2D complex conv-block with kernel size of 5x4. RAKI is a CNN with 3 conv-blocks whose kernel sizes are 5x2, 1x1, 3x2 with ReLU activation functions. Original RAKI implemented the convolution by concatenating the real and imaginary components. Here, we also explore the use of complexRAKI (cRAKI) which exploits a true complex implementation: the convolution is replaced by complex convolution and LeakyReLU is replaced with complex LeakyReLU [17].

Patch discriminator. While GRAPPA is a standard reconstruction method that has been imple-

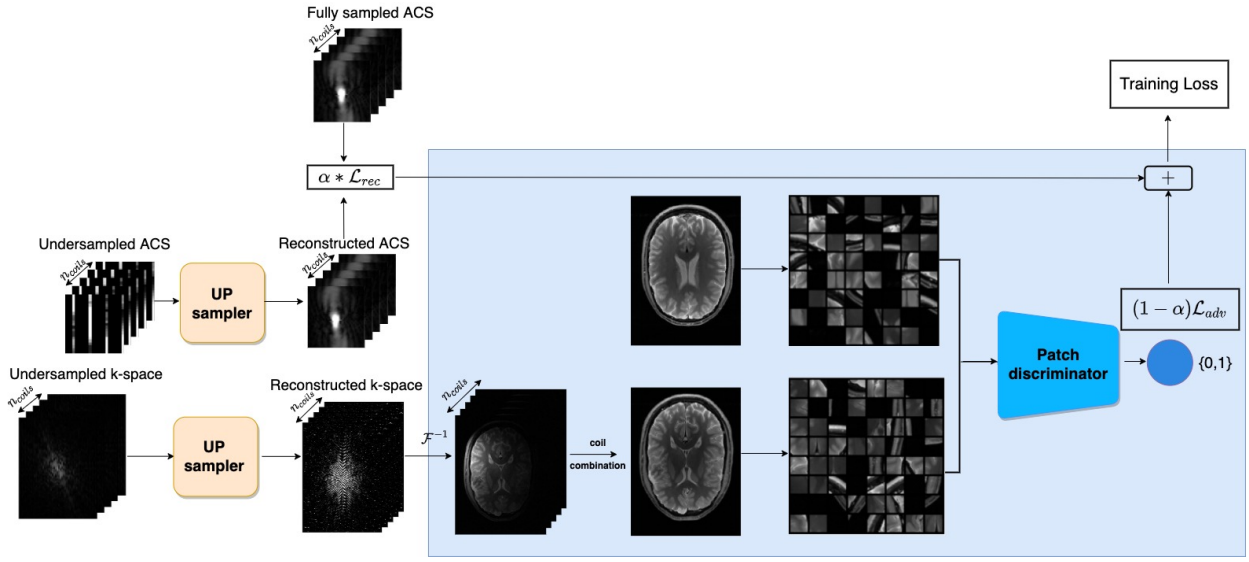

Figure 2: The architecture of the proposed one-shot learning Super-ResolutIon with Adversarial training , SANGRIA. It consists of super-resolution model (upsampler) coupled with a patch-discriminator. Super-resolution models are either GRAPPA, RAKI or cRAKI.

mented on MRI systems for many years, the reconstucted image quality quickly degrades with the acceleration rate which is typically 2-3 in clinical practice. In 2D imaging, R above 3 usually yields unusable images for diagnosis.

We therefore introduce a discriminator to increase the quality of the upsampled image as it has been proposed for natural image super-resolution [13] and more recently for MRI images [14]. As the goal is to learn a super-resolution model for each slice, one can only rely on a patch-based adversarial discriminator, which has to be learned on a few data. To cope with this particular setting, we experimented with a shallow patch GAN discriminator with limited receptive field. We choose it to accept input as 17x17 patch. The discriminator consists of three convolution blocks, each including a convolution layer, batch normalization and LeakyReLU, followed by a fully connected layer with a single decision output. Convolution layers have successively 1, 32 and 64 filters with size equal to 5×5. The discriminator takes the coil-combined image as input, and outputs a decision score for each patch.

Training. The full model is trained with a combined loss which is the sum of the reconstruction loss and of the adversarial loss.

$$\min_G \max_D \alpha \mathcal{L}_{rec}(G) + (1-\alpha)\mathcal{L}_{adv}(G,D) \qquad (5)$$

where α is a regularization parameter bewteen 0 and 1 that balances the two losses. The reconstruction loss is the mean squared error loss between the reconstructed ACS $\hat{y}$ and the fully-sampled ACS y, restricted to the training partition.

$$\mathcal{L}_{rec}(y,\hat{y}) = \frac{1}{N}\sum_{j=1}^{N}(y_j - \hat{y}_j)^2 \qquad (6)$$

where N denotes the total number of k-space points. The adversarial loss $\mathcal{L}_{adv}$ penalizes the reconstructor for not generating realistic patches. D is trained to classify correctly the patches coming from both the patches from the fully sampled image and from the reconstucted image generated by G. The fully sampled image we call x_{ref} is the single fully sampled coil-combined (SoS) in the image space. Simultaneously G is trained to fool the discriminator and to generate a reconstructed k-space $G(S)$ (with S the multi-coils undersampled k-space), whose patches (after transformation in the image space) are classified as realistic by the discriminator.

$$\mathcal{L}_{adv}(G,D)(S) = E_{p\sim x_{ref}}[\log D(p)] + E_{p\sim i_R}[1 - \log D(p)] \qquad (7)$$

where i_R stands for the reconstructed image in the image space $i_R = \sqrt{\sum_{i=1}^{n_c}(\mathcal{F}^{-1}G(S_i))^2}$ from the inputs coils S, $p \sim x_{ref}$ means that p follows the patch distribution of x_{ref} and $p \sim i_R$ means that p follows the patch distribution of i_R.

5 Experiments

5.1 Data and experimental setting

Data. We performed experiments on both synthetic brain images and in-vivo brain data acquired from a healthy volunteer. In-vivo data consists of various multi-slice T2-weighted 2D spin-echo brain acquisitions with in-plane resolution of 0.8 mm, matrix size 256x196 at 3 Tesla (Siemens Vida). Informed consent was obtained. Fully-sampled data (acquisition time T_{acq} of 4 minutes), as well as under-sampled data with acceleration factors ranging from R=2 to 6, were acquired, resulting in R times faster acquisitions. ACS of size (49,52,64,40) is acquired in the separated fast scan. Concerning synthetic data, brain volumes from the MNI152 standard-space T1-weighted average structural template were generated [18]. The template provides tissue probability maps of the following tissue classes: grey matter, white matter, cerebrospinal fluid, bone, skin and air. The MNI template is provided in the Statistical Parametric Mapping (SPM12) Matlab software. Synthetic brain volumes were simulated by setting a different signal level for each class. Coil sensitivities were simulated using Biot and Savart law assuming an $n_c = 8$ receiver-coil array with coils positioned regularly spaced around the brain on a cylindrical surface. From these simulated coil sensitivities and synthetic brain volumes, fully-sampled and under-sampled k-spaces with different acceleration factor R and noise levels, as well as SoS-image combinations were generated.

Experimental setting. All the methods were trained for a maximum of 1000 epochs with early stopping. The reconstruction model and the discriminator are learned with learning rates that are grid searched on set by cross-validation (learning rates from 1e-4 to 1e-2). In any case, training ends when performance on the validation dataset starts to degrade. We used Adam optimizer with $\beta 1 = 0.9$, $\beta 2 = 0.999$.

Metrics. Qualitative and quantitative evaluations for comparison of the various methods were performed. For quantitative evaluation, Normalized Root Mean Squared Error (NRMSE), Normalized Root Mean Absolute Error (NMAE), Peak Signal to Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM) [19] were calculated. For qualitative evaluation, examples of reconstructed images were inspected visually.

5.2 Quantitative evaluation of the reconstruction quality

In the first series of experiments, we compare all the methods on data where train and test images correspond to the same acceleration rate, and uniform undersampling mask. We provide results for simulated data and for in-vivo data for uniform undersampling.

Synthetic data Table 1 shows the quantitative comparison of various methods for the simulated data. We observe that at $R = 4, 5$ cRAKI has superior performance in terms of SSIM and PSNR as compared to GRAPPA and RAKI. When GRAPPA is coupled with an adversarial discriminator D, we observe improved performance for both $R = 4, 5$ and for all the metrics. Besides the addition of the adversarial discriminator in RAKI and cRAKI modeling sighlty degrades the performance in terms of SSIM/PSNR but sighlty improves with respect to NRMSE/MAE both for $R = 4, 5$.

In-vivo data Table 2 compares the same quantitative metrics of various methods as above but for volunteer data, averaged over 49 slices, and with a higher acceleration rate of 5 and 6. A larger acceleration rate could be obtained in vivo. In this case, data were acquired with n_c=52 receiver coils (while the synthetic data had n_c=8), and acquisition field-of-view was slightly larger than the brain dimension along the phase encoding. This could explain that the acceleration rate that can be achieved in vivo is larger than the one reached in simulation. Unlike the simulated data, image qualitative metrics such as SSIM/PSNR decreased when GRAPPA and RAKI is coupled with discriminator at both R=5,6. However, we observed an improvement in NRMSE/NMAE when RAKI was trained with the discriminator at both R=5,6. cRAKI and cRAKI with D has similar quantitative metrics at R=5, but it degrades at R=6.

5.3 Qualitative evaluation of reconstruction quality

While prior results have not clearly demonstrated an advantage to using an adversarial discriminator with respect to objective evaluation measures, including image quality-based criteria such as SSIM/PSNR, we provide here qualitative results that demonstrate that these objective criteria do not directly relate to subjective image quality. It is been reported in the computer vision literature [13] that perceptual quality may not corresponds to the objective metrics.

Reconstructor	**R**	**NRMSE**	**NMAE**	**SSIM**	**PSNR**
GRAPPA	4x	0.058±0.060	0.191±0.088	0.943±0.080	38.322±4.463
RAKI	4x	0.066±0.082	0.215±0.084	0.971±0.013	42.785±6.845
cRAKI	4x	0.060±0.085	0.164±0.105	0.984±0.009	44.599±6.267
GRAPPA w/ D	4x	0.050±0.017	0.154±0.028	0.982±0.011	39.526±1.606
RAKI w/ D	4x	0.063±0.064	0.216±0.069	0.942±0.041	37.440±2.886
cRAKI w/ D	4x	0.025±0.015	0.112±0.055	0.974±0.029	44.081±2.680
GRAPPA	5x	0.217±0.239	0.378±0.147	0.808±0.189	28.174±8.813
RAKI	5x	0.541±0.138	0.523±0.074	0.680±0.096	23.893±9.801
cRAKI	5x	0.083±0.076	0.245±0.079	0.966±0.014	40.769±7.360
GRAPPA w/ D	5x	0.113±0.032	0.380±0.049	0.895±0.023	30.551±2.232
RAKI w/ D	5x	0.181±0.250	0.368±.159	0.799±0.221	30.700±8.094
cRAKI w/ D	5x	0.047±0.032	0.188±0.061	0.958±0.031	39.100±1.892

Table 1: Average metrics over 121 slices for the synthetic data for various reconstructors coupled (or not) with an adversarial patch-discriminator at R=4,5. We performed experiments on uniform undersampling.

Reconstructor	**R**	**NRMSE**	**NMAE**	**SSIM**	**PSNR**
GRAPPA	5x	0.182±0.046	0.545±0.087	0.946±0.016	37.176±1.665
RAKI	5x	0.209±0.045	0.514±0.072	0.955±0.114	39.276±2.030
cRAKI	5x	0.158±0.038	0.461±0.066	0.898±0.033	36.401±1.835
GRAPPA w/ D	5x	0.239±0.082	0.725±0.173	0.902±0.035	35.013±2.157
RAKI w/ D	5x	0.164±0.039	0.482±0.069	0.920±0.022	36.755±1.668
cRAKI w/ D	5x	0.158±0.038	0.463±0.066	0.901±0.032	36.451±1.809
GRAPPA	6x	0.354±0.046	0.769±0.074	0.844±0.054	28.846±2.539
RAKI	6x	0.265±0.044	0.606±0.063	0.915±0.022	36.088±2.099
cRAKI	6x	0.215±0.046	0.576±0.067	0.867±0.032	33.099±1.701
GRAPPA w/ D	6x	0.427±0.097	1.102±0.190	0.732±0.055	27.079±1.982
RAKI w/ D	6x	0.230±0.055	0.647±0.090	0.876±0.032	33.182±1.926
cRAKI w/ D	6x	0.268±0.072	0.703±0.105	0.854±0.058	32.361±2.532

Table 2: Average metrics over 49 slices for the volunteer data for various reconstructors coupled (or not) with an adversarial patch-discriminator, D at R=5,6. The MRI zero-shot methods such as GRAPPA, RAKI and cRAKI is trained with/without adversarial tranining. We performed experiments on uniform undersampling.

The first series of images in figure 3 provides a qualitative comparison of the various methods. It show superior visual quality of reconstucted images when the methods are coupled with D. RAKI perform better in terms of the image-quality index such as SSIM/PSNR. However, we observed the images obtained from RAKI has more undersampling artefacts than GRAPPA and artefactual 'dots' around the centre region. When RAKI is coupled with D, artefactual 'dots' disappears and SSIM/PSNR is degraded. NRMSE/NMAE is lower in cRAKI and cRAKI with D, indicating that the model better represents the ACS data. However, in terms of image-based quantitative measured SSIM/PSNR degrades a little, although Figure 3 indicates cRAKI with D improved visual quality both R=5,6. The cRAKI reconstruction has fewer undersampling artefacts than GRAPPA and the absence of artefactual 'dots' (red square) appeared in RAKI reconstruction. RAKI and cRAKI with D both can perform reconstruction up to R=6. Although, we obtained noisy reconstructions at R=6 but the global features of the brain image were restored, while all other methods failed to reconstruct at this acceleration rate.

The second series of images in figure 4 illustrates how the learning changes when the tradeoff between the reconstruction loss and the adversarial loss is varied (through the α hyperparameter in Eq.5. Various values of α are used in the range of 1e-6 to 1. The quantitative measures are similar for all the α greater than 1e-3. Visually, however, the reconstruction quality improves for $\alpha = 0.1$. When $\alpha = 1$, only reconstruction loss dominates, and undersampling artefacts are visible.

6 Conclusion

We proposed a one-shot learning approach for MRI reconstruction, which combines state-of-the-art zero-shot MRI super-resolution models operating in kspace with an adversarial patch discriminator operating in the image domain. We present a train-validation partition that is specific to uniform undersampling in order to avoid overfitting during training. The subjective analysis shows improved-quality reconstruction when an adversarial discriminator is introduced. Due to the lack of an objective criterion which strongly related to the subjective quality of the reconstructed image, a large-scale subjective evaluation could be performed in order to confirm the superiority of this adversarial strategy for MRI reconstruction.

Acknowledgement

This work was supported by the French National Research Agency (ANR) - Doctoral contracts in Artificial Intelligence, Archimède Institute, Institute Marseille Imaging, A*MIDEX excellence initiative.

References

[1] P. B. Roemer, W. A. Edelstein, C. E. Hayes, S. P. Souza, and O. M. Mueller, "The NMR phased array," en, *Magnetic Resonance in Medicine*, vol. 16, no. 2, pp. 192–225, 1990.

[2] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, "SENSE: Sensitivity encoding for fast MRI," eng, *Magnetic Resonance in Medicine*, vol. 42, no. 5, pp. 952–962, Nov. 1999.

[3] M. A. Griswold, P. M. Jakob, R. M. Heidemann, *et al.*, "Generalized autocalibrating partially parallel acquisitions (GRAPPA)," *Magnetic Resonance in Medicine*, vol. 47, no. 6, pp. 1202–1210, Jun. 2002.

[4] A. Sriram, J. Zbontar, T. Murrell, *et al.*, "End-to-end variational networks for accelerated MRI reconstruction," in *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*, Cham: Springer International Publishing, 2020, pp. 64–73.

[5] B. Zhu, J. Z. Liu, S. F. Cauley, B. R. Rosen, and M. S. Rosen, "Image reconstruction by domain-transform manifold learning," en, *Nature*, vol. 555, no. 7697, pp. 487–492, Mar. 2018.

[6] M. Akçakaya, S. Moeller, S. Weingärtner, and K. Uğurbil, "Scan-specific robust artificial-neural-networks for k-space interpolation (RAKI) reconstruction: Database-free deep learning for fast imaging," *Magnetic Resonance in Medicine*, vol. 81, no. 1, pp. 439–453, 2019.

[7] B. Yaman, S. A. H. Hosseini, and M. Akcakaya, "Zero-shot self-supervised learning for MRI reconstruction," presented at the International Conference on Learning Representations, Sep. 29, 2021.

[8] K. Hammernik, T. Klatzer, E. Kobler, *et al.*, "Learning a variational network for reconstruction of accelerated MRI data," *Magnetic Resonance in Medicine*, vol. 79, no. 6, pp. 3055–3071, 2018.

[9] G. Yang, S. Yu, H. Dong, *et al.*, "DAGAN: Deep De-Aliasing Generative Adversarial Networks for Fast Compressed Sensing MRI Reconstruction," *IEEE Transactions on Medical Imaging*, vol. 37, no. 6, pp. 1310–1321, Jun. 2018.

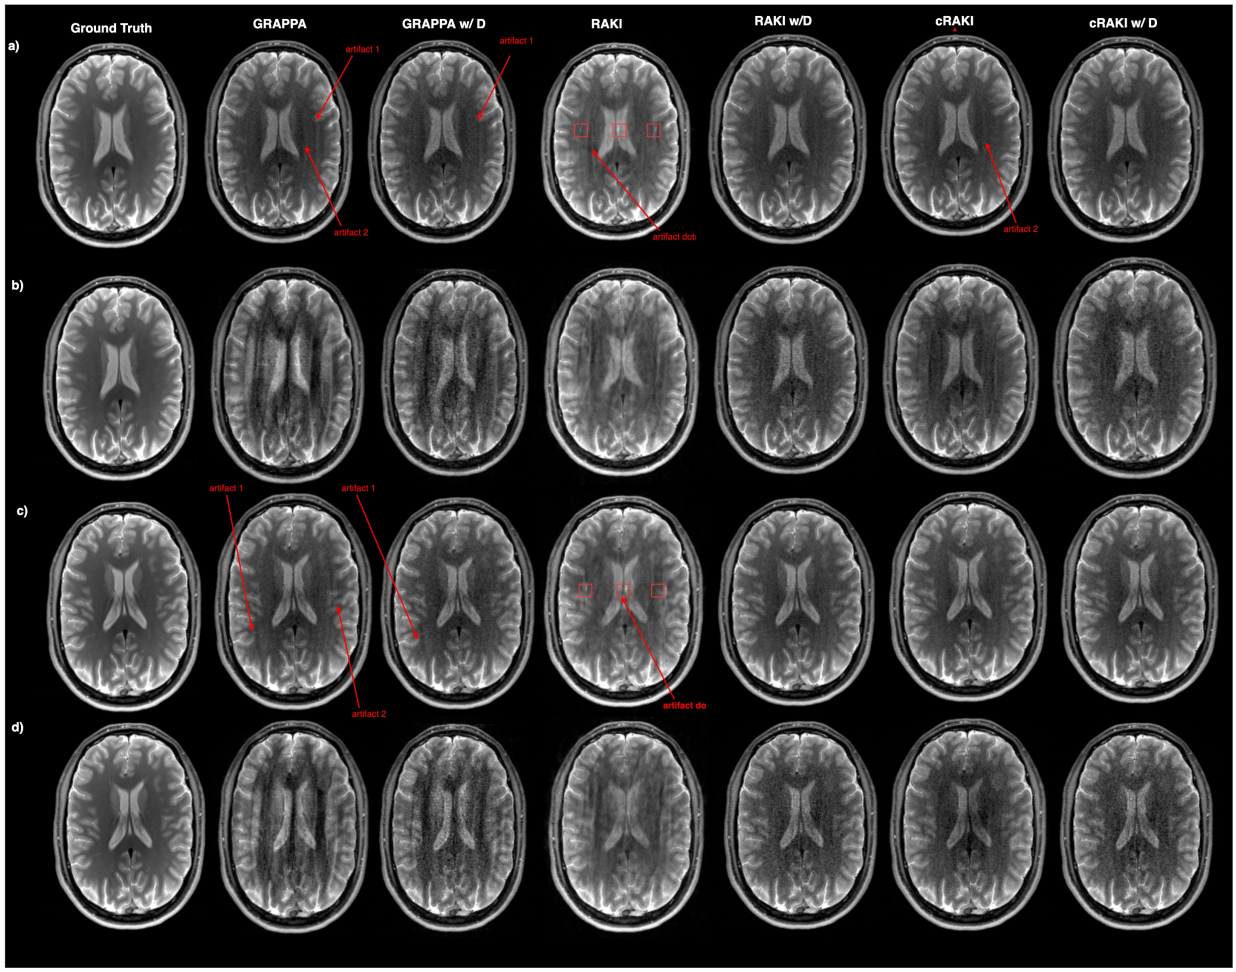

Figure 3: Reconstruction results on various slices where the test slice corresponds to the same contrast, acceleration rate and undersampling mask as the train slice for a) R=5 at 1st slice b) R=6 at 1st slice c) R=5 at 2nd slice d) R=5 at 2nd slice. The red arrows represent the aliasing artefacts from the undersampling. GRAPPA, GRAPPA w/ D, RAKI cRAKI have the undersampling . RAKI reconstruction contains some artefactual dots at the centre represented in red squares. RAKI w/D do not display these artefactual dots. cRAKI w/ D and RAKI w/ D both have clinically acceptable image quality at R=5. At R=6, all the methods fail to reconstruct except RAKI and cRAKI w/ D.

[10] T. M. Quan, T. Nguyen-Duc, and W.-K. Jeong, "Compressed Sensing MRI Reconstruction Using a Generative Adversarial Network With a Cyclic Loss," *IEEE Transactions on Medical Imaging*, vol. 37, no. 6, pp. 1488–1497, Jun. 2018.

[11] M. Bento, I. Fantini, J. Park, L. Rittner, and R. Frayne, "Deep Learning in Large and Multi-Site Structural Brain MR Imaging Datasets," *Frontiers in Neuroinformatics*, vol. 15, 2022.

[12] R. V. Zicari, J. Brusseau, S. N. Blomberg, *et al.*, "On Assessing Trustworthy AI in Healthcare. Machine Learning as a Supportive Tool to Recognize Cardiac Arrest in Emergency Calls," *Frontiers in Human Dynamics*, vol. 3, 2021.

[13] C. Ledig, L. Theis, F. Huszar, *et al.*, "Photo-realistic single image super-resolution using a generative adversarial network," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, Jul. 2017, pp. 105–114.

[14] J. Wang, Y. Chen, Y. Wu, J. Shi, and J. Gee, "Enhanced generative adversarial network for 3d brain MRI super-resolution," in *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*, IEEE, Mar. 2020, pp. 3616–3625.

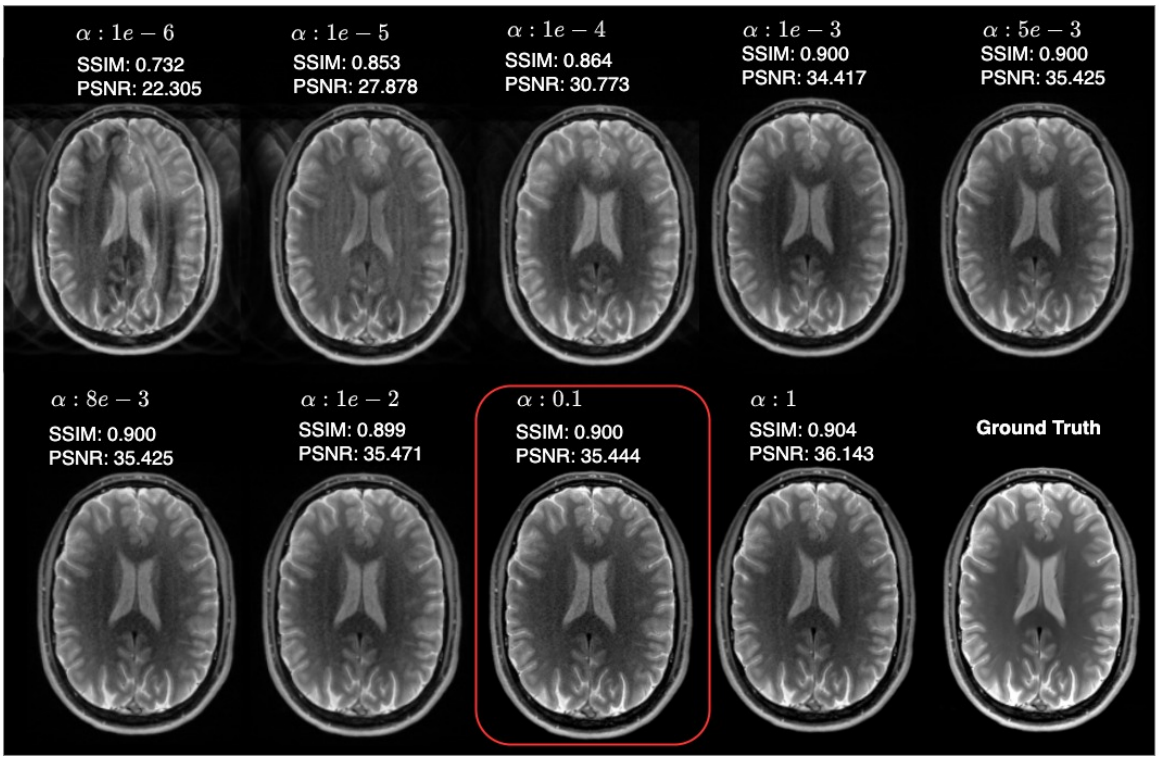

Figure 4: Effect of α in Eq.5, trade-off parameter between the reconstruction loss and the adversarial loss. The reconstructions is performed with cRAKI at R=5 and varying α from 1e-6 to 1. Generally, the lower the value of α, the smaller the contribution of reconstruction loss

[15] S. Iizuka, E. Simo-Serra, and H. Ishikawa, "Globally and Locally Consistent Image Completion," English, *Acm Transactions on Graphics*, vol. 36, no. 4, p. 107, Jul. 2017.

[16] C. Zhang, S. A. H. Hosseini, S. Weingärtner, K. Uğurbil, S. Moeller, and M. Akçakaya, "Optimized fast GPU implementation of robust artificial-neural-networks for k-space interpolation (RAKI) reconstruction," *PLOS ONE*, vol. 14, no. 10, Oct. 23, 2019.

[17] C. Trabelsi, O. Bilaniuk, Y. Zhang, *et al.*, "Deep complex networks," presented at the International Conference on Learning Representations, Feb. 15, 2018.

[18] V. Fonov, A. Evans, R. McKinstry, C. Almli, and D. Collins, "Unbiased nonlinear average age-appropriate brain templates from birth to adulthood," *NeuroImage*, Organization for Human Brain Mapping 2009 Annual Meeting, vol. 47, S102, Jul. 1, 2009.

[19] J. Zbontar, F. Knoll, A. Sriram, *et al.*, "fastMRI: An open dataset and benchmarks for accelerated MRI," *arXiv:1811.08839 [physics, stat]*, Dec. 11, 2019.