

DAGAN: Deep De-Aliasing Generative Adversarial Networks for Fast Compressed Sensing MRI Reconstruction

Guang Yang*, Member, IEEE, Simiao Yu*, Hao Dong, Greg Slabaugh, Senior Member, IEEE, Pier Luigi Dragotti, Fellow, IEEE, Xujiong Ye, Fangde Liu, Simon Arridge, Jennifer Keegan, Yike Guo**, and David Firmin**

Abstract—Compressed Sensing Magnetic Resonance Imaging (CS-MRI) enables fast acquisition, which is highly desirable for numerous clinical applications. This can not only reduce the scanning cost and ease patient burden, but also potentially reduce motion artefacts and the effect of contrast washout, thus yielding better image quality. Different from parallel imaging based fast MRI, which utilises multiple coils to simultaneously receive MR signals, CS-MRI breaks the Nyquist-Shannon sampling barrier to reconstruct MRI images with much less required raw data. This paper provides a deep learning based strategy for reconstruction of CS-MRI, and bridges a substantial gap between conventional non-learning methods working only on data from a single image, and prior knowledge from large training datasets. In particular, a novel conditional Generative Adversarial Networks-based model (DAGAN) is proposed to reconstruct CS-MRI. In our DAGAN architecture, we have designed a refinement learning method to stabilise our U-Net based generator, which provides an end-to-end network to reduce aliasing artefacts. To better preserve texture and edges in the reconstruction, we have coupled the adversarial loss with an innovative content loss. In addition, we incorporate frequency domain information to enforce similarity in both the image and frequency domains. We have performed comprehensive comparison studies with both conventional CS-MRI reconstruction methods and newly investigated deep learning approaches. Compared to these methods, our DAGAN method provides superior reconstruction with preserved perceptual image details. Furthermore, each image is reconstructed in about 5 ms, which is suitable for real-time processing.

Index Terms—Compressed Sensing; Magnetic Resonance Imaging (MRI); Fast MRI; Deep Learning; Generative Adversarial Networks (GAN); De-aliasing; Inverse Problems.

G. Yang is with the National Heart & Lung Institute, Imperial College London, SW3 6NP, London, United Kingdom. g.yang@imperial.ac.uk

S. M. Yu, H. Dong, F. D. Liu and Y. K. Guo are with the Data Science Institute, Imperial College London, SW7 2AZ, London, United Kingdom. {simiao.yu13, hao.dong11, fangde.liu, y.guo}@imperial.ac.uk

G. Slabaugh is with the Dept. of Computer Science, City University of London, EC1V 0HB, London, United Kingdom. Gregory.Slabaugh.1@city.ac.uk

P. L. Dragotti is with the EEE Department, Imperial College London, SW3 6NP, London, United Kingdom. p.dragotti@imperial.ac.uk

X. J. Ye is with the School of Computer Science, University of Lincoln, LN6 7TS, Lincoln, United Kingdom. xye@lincoln.ac.uk

S. Arridge is with the CMIC, University College London, WC1E 6BT, London, United Kingdom. s.arridge@cs.ucl.ac.uk

J. Keegan and D. Firmin are with the National Heart & Lung Institute, Imperial College London, SW3 6NP, London, United Kingdom. {j.keegan,d.firmin}@imperial.ac.uk

G. Yang, J. Keegan and D. Firmin are supported by the British Heart Foundation Project Grant (Project Number: PG/16/78/32402), the NIHR Cardiovascular Biomedical Research Unit, Royal Brompton Hospital & Harefield NHS Foundation Trust and Imperial College London; S. M. Yu is funded by the Jaywing plc and H. Dong is funded by the OPTIMISE Portal.

Asterisk indicates co-first authors, Y. K. Guo and D. Firmin are co-last senior authors.

I. INTRODUCTION

MAGNETIC Resonance Imaging (MRI) is a widely applied medical imaging modality for numerous clinical applications. MRI can provide reproducible, non-invasive, and quantitative measurements of tissue, including structural, anatomical and functional information. However, one major drawback of MRI is the prolonged acquisition time. MRI is associated with an inherently slow acquisition speed that is due to data samples not being collected directly in the image space but rather in k -space. K -space contains spatial-frequency information that is acquired line-by-line and anywhere from 64 to 512 lines of data are needed for a high quality reconstruction. This relatively slow acquisition could result in significant artefacts due to patient movement and physiological motion, e.g., cardiac pulsation, respiratory excursion, and gastrointestinal peristalsis. Prolonged acquisition times also limit the usage of MRI due to expensive cost and considerations of patient comfort and compliance [1]. Moreover, for protocols that require contrast agent injection, lengthy acquisition can result in contrast washout that may lead to poor quality or non-diagnostic images. Due to limitations of the scanning speed, patient throughput using MRI is slow compared with other medical imaging modalities.

The MRI raw data samples are acquired sequentially in k -space and the speed at which k -space can be traversed is limited by physiological and hardware constraints [2]. Once the desired field-of-view and spatial resolution of the MRI images are prescribed, the required k -space raw data is conventionally determined by the Nyquist-Shannon sampling criteria [3]. Some early research on fast MRI proposed acquiring several lines in k -space from a single radio frequency (RF) excitation by implementing multiple RF [4] or gradient [5] refocussings. Since these acceleration techniques acquire the full k -space coverage demanded by the Nyquist-Shannon sampling criteria, they are still categorised as *fully sampled* methods [1].

One possible fast MRI approach is to undersample k -space, resulting in an acceleration rate proportional to the undersampling ratio. Partial Fourier imaging (PFI) [6] is an *undersampled* technique based on the principle that in theory, only half of the k -space in the phase encoding direction is required according to the property that the Fourier transformation of a purely real function has complex conjugate symmetry in k -space [7]. However, in practice, more than half of the phase encoding is acquired to provide a robust phase correction [1]; therefore, the acceleration factor using PFI is limited to <2 and it is associated with a drop of signal to noise ratio (SNR). Alternatively, parallel imaging is a fast MRI method

using multiple independent receiver channels. Each independent channel is most sensitive to the tissue nearest to that coil. The raw data acquired from these independent channels can be combined using either a sensitivity encoding (SENSE) technique [8] or a generalised auto-calibrating (GRAPPA) method [9]. The acceleration factor of parallel imaging is limited by the number and arrangement of the receiver coils, which potentially introduce some imaging artefacts [10] and increase the manufacturing cost of the MRI scanner.

On the other hand, Compressed Sensing [11] based MRI (CS-MRI) allows fast acquisition that bypasses the Nyquist-Shannon sampling criteria with more aggressive undersampling. In theory, it can achieve a reconstruction without deterioration of image quality by performing *nonlinear optimisation* on *randomly undersampled* raw data, assuming the data is *compressible*. The main challenge for CS-MRI is to find an algorithm that can reconstruct an uncorrupted or de-aliased image from randomly highly undersampled k -space data.

II. RELATED WORK AND OUR CONTRIBUTIONS

A. Classic Model-Based CS-MRI

CS-MRI research has been focused on three major directions. First, investigations have sought the best undersampling scheme, which should be as random as possible to create incoherent undersampling artefacts so that a proper nonlinear reconstruction can be applied to suppress noise-like artefacts without degrading image quality of the reconstruction [12], [13], [14]. The sparsity of the random undersampling determines the acceleration rate. More importantly, the random undersampling scheme must be feasible to implement on the MRI scanner and compatible with particular scanning sequences. Currently, most studies using 2D MRI have relied on 1D random undersampling to produce a sampling pattern that follows a 1D Gaussian distribution. This gives a higher level of undersampling in the high frequency regions while low frequencies are retained to preserve overall image structure. Only 1D undersampling is used because the sampling along the frequency encoding direction is fast, and only the phase encoding direction limits the acquisition time [2]. Considering 3D reconstruction, 2D Gaussian and Poisson disc masks are commonly used to accelerate phase and slice encoding [1].

Second, in general, medical imagery acquired by MRI is naturally compressible. CS-MRI utilises the implicit sparsity to reconstruct accelerated acquisitions [15]. Here the term sparsity describes a matrix of image pixels or raw data points which are predominately zero valued or namely compressible. Such sparseness may exist either in the image domain or more commonly via a suitable mathematical representation in a transform domain. Sparse representation can be explored in a specific transform domain or generally in a dictionary-based subspace [16]. Classic fast CS-MRI uses predefined and fixed sparsifying transforms, e.g., total variation (TV) [17], [18], [19], discrete cosine transforms [20], [21], [22] and discrete wavelet transforms [23], [24], [25]. In addition, this has been extended to a more flexible sparse representation learnt directly from data using dictionary learning [26], [27], [28].

Finally, nonlinear optimisation algorithms ensure the achievement of efficient, stable and accurate reconstruction

[16], [18], [29], [30]. Comprehensive reviews on classic CS-MRI methods and clinical applications can be found elsewhere, e.g., [1], [31].

Although there are promising studies applying fast CS-MRI in clinical environments [31], [32], [33], most routine clinical MRI scanning is still based on standard fully-sampled Cartesian sequences or is accelerated only using parallel imaging. The main challenges are: (1) satisfying the incoherence criteria required by CS-MRI [1]; (2) the widely applied sparsifying transforms might be too simple to capture complex image details associated with subtle differences of biological tissues, e.g., TV based sparsifying transform penalises local variation in the reconstructed images but can introduce staircase artefacts and the wavelet transform enforces point singularities and isotropic features but orthogonal wavelets may lead to blocky artefacts [34], [35], [36]; (3) nonlinear optimisation solvers usually involve iterative computation that may result in relatively long reconstruction time [1]; (4) inappropriate hyperparameters predicted in current CS-MRI methods can cause over-regularisation that will yield overly smooth and unnatural looking reconstructions or images with residual undersampling artefacts [1]. Due to these challenges and limitations, the acceleration rate using CS-MRI alone is still limited ($2\times$ to $6\times$ acceleration).

B. Deep Learning-Based CS-MRI

Recently, deep learning has received great attention in computer vision studies and has generally returned dividends in performance. Shen et al. [37] surveyed the most recent research in deep learning for medical image analysis and Wang [38] provided an insightful perspective on deep imaging that proposed to incorporate deep learning into tomographic image reconstruction. Essentially, CS-MRI reconstruction solves a generalised inverse problem that is analogous to image super-resolution (SR) [39], de-noising and inpainting [40], [41] that have been successfully solved using deep neural network architectures, e.g., using convolutional neural networks (CNN).

Currently, there are only preliminary studies on deep learning based CS-MRI reconstruction. Wang et al. [42] introduced a CNN-based CS-MRI method, in which the learnt network was used to initialise the classic CS-MRI in a two-phase reconstruction, or integrated into the CS-MRI directly as an additional regularisation term. Despite preliminary qualitative visualisation that showed some promise, the applicability of this method for CS-MRI reconstruction is yet to be quantitatively assessed in detail. In [36], a deep network was trained for solving CS-MRI using an Alternating Direction Method of Multipliers [18] framework, i.e., ADMM-Net. The reconstruction, de-noising and Lagrangian multiplier updates were implemented in a data flow graph and optimised through cascaded deep network layers. The method achieved similar reconstruction results as classic CS-MRI reconstruction methods but dramatically reduced reconstruction time. There are also three preprints describing deep learning based CS-MRI: Schlemper et al. [43] proposed cascaded CNN incorporating a data consistency layer, Hammernik et al. [44] trained a variational network to solve CS-MRI and Lee et

al. [45] combined CNN based CS-MRI with parallel imaging to estimate and remove the aliasing artefacts. Although deep learning has shown great potential in solving CS-MRI with much faster reconstruction, to date improvement was not found significantly different from what classic CS-MRI can achieve. Moreover, similar to other deep learning applications, it is not trivial to define the network architecture and convergence of the deep network training might be difficult to achieve unless comprehensive parameter tuning is performed.

C. Our Contributions

In this study, we proposed a novel conditional Generative Adversarial Networks (GAN) based deep learning architecture (dubbed DAGAN) for de-aliasing and fast CS-MRI by a comprehensive extension of our initial proof-of-concept study [46] in both method and simulation settings. Our main contributions are:

- We propose a U-Net architecture [47], [48] with skip connections for the generator network;
- A refinement learning approach is designed to stabilise the training of GAN for fast convergence and less parameter tuning;
- The adversarial loss is coupled with a novel content loss considering both pixel-wise mean square error (MSE) and perceptual loss defined by pretrained deep convolutional networks from the Visual Geometry Group at Oxford University (in short VGG networks [49]) to achieve better reconstruction details;
- Frequency domain information of the CS-MRI has been incorporated as additional constraints for the data consistency, which is formed as an extra loss term;
- We perform comprehensive experiments and compare our proposed models with both classic CS-MRI and newly developed deep learning based methods.

Compared to the state-of-the-art CS-MRI methods, we can achieve high acceleration factors with superior results and faster processing time.

III. METHOD

A. General CS-MRI

Forward Model The observation or data acquisition forward model of image restoration or reconstruction can be approximated as a discretised linear system [50],

$$\mathbf{F}\mathbf{x} + \varepsilon = \mathbf{y}, \quad (1)$$

in which $\mathbf{x} \in \mathbb{C}^N$ represents the desired image to be reconstructed, which consists of $\sqrt{N} \times \sqrt{N}$ pixels formatted as a column vector. The observation is denoted as $\mathbf{y} \in \mathbb{C}^M$. The forward model of image acquisition can be described using a linear operator $\mathbf{F} : \mathbb{C}^N \mapsto \mathbb{C}^M$ that different matrices $\mathbf{F} \in \mathbb{C}^{M \times N}$ represent various image restoration or reconstruction problems, e.g., an identity operator for image de-noising, convolution operators for de-blurring, filtered subsampling operators for SR and k -space random undersampling operators for CS-MRI reconstruction [50].

Inverse Model The inverse estimation of Eq. 1 is usually ill-posed because the problem is normally underdetermined

with $M \ll N$. Moreover, the inverse model is unstable due to a numerically ill-conditioned operator \mathbf{F} and the presence of noise (ε in Eq. 1) [50].

Classic Model-Based CS-MRI In order to solve this underdetermined and ill-posed system of CS-MRI, one must exploit a-priori knowledge of \mathbf{x} that can be formulated as an unconstrained optimisation problem, that is

$$\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{F}_u \mathbf{x} - \mathbf{y}\|_2^2 + \lambda \mathcal{R}(\mathbf{x}), \quad (2)$$

where $\frac{1}{2} \|\mathbf{F}_u \mathbf{x} - \mathbf{y}\|_2^2$ is the data fidelity term and $\mathbf{F}_u \in \mathbb{C}^{M \times N}$ is the undersampled Fourier encoding matrix. \mathcal{R} expresses regularisation terms on \mathbf{x} and λ is a regularisation parameter. The regularisation terms \mathcal{R} typically involve l_q -norms ($0 \leq q \leq 1$) in the sparsifying domain of \mathbf{x} [2].

Deep Learning-Based CS-MRI Deep learning based studies [42], [43] propose to incorporate a CNN into CS-MRI reconstruction, that is

$$\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{F}_u \mathbf{x} - \mathbf{y}\|_2^2 + \lambda \mathcal{R}(\mathbf{x}) + \zeta \|\mathbf{x} - f_{\text{cnn}}(\mathbf{x}_u | \hat{\theta})\|_2^2, \quad (3)$$

in which f_{cnn} is the forward propagation of data through the CNN parametrised by θ , and ζ is another regularisation parameter. The image generated by the CNN (i.e., $f_{\text{cnn}}(\mathbf{x}_u | \hat{\theta})$) is used as a reference image and as an additional regularisation term, in which $\hat{\theta}$ represents the optimised parameters of the trained CNN. In addition, $\mathbf{x}_u = \mathbf{F}_u^H \mathbf{y}$ is the reconstruction from the zero-filled undersampled k -space measurements, where H represents the Hermitian transpose operation. MRI data naturally encodes magnitude and phase information in complex number format. There are at least two strategies for a deep learning based method to handle complex numbers: (1) real-valued information can be embedded into the complex space using an operator $\text{Re}^* : \mathbb{R}^N \mapsto \mathbb{C}^N$ such that $\text{Re}^*(\mathbf{x}) = \mathbf{x} + 0i$, and therefore the MRI forward operator can be expressed as $\mathbf{F} : \mathbb{R}^N \xrightarrow{\text{Re}^*} \mathbb{C}^N \xrightarrow{\mathcal{F}} \mathbb{C}^N \xrightarrow{\mathcal{U}} \mathbb{C}^M$ that \mathbf{F}_u combines the Fourier transform \mathcal{F} and random undersampling \mathcal{U} operators [51]; (2) the real and imaginary data can be learnt as separate channels in the training of the CNN, in other words, the \mathbb{C}^N is replaced as \mathbb{R}^{2N} [43]. In our simulation based study, the first strategy was employed to avoid extra computational burden.

B. General GAN

Generative Adversarial Networks [52] consist of a generator network G and a discriminator network D . The goal of the generator G is to map a latent variable \mathbf{z} , e.g., an input vector of random numbers, to the distribution of the given true data \mathbf{x} that we are interested in imitating in order to fool the discriminator D . The discriminator aims to distinguish the true data \mathbf{x} from the synthesized fake data $G_{\theta_G}(\mathbf{z})$. GAN can be formulated mathematically as a minimax game between the generator $G_{\theta_G}(\mathbf{z}) : \mathbf{z} \mapsto \mathbf{x}$ and the discriminator $D_{\theta_D}(\mathbf{x}) : \mathbf{x} \mapsto [0, 1]$, and this training process is parameterised by θ_G and θ_D as following

$$\min_{\theta_G} \max_{\theta_D} \mathcal{L}(\theta_D, \theta_G) = \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}(\mathbf{x})} [\log D_{\theta_D}(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [\log(1 - D_{\theta_D}(G_{\theta_G}(\mathbf{z})))] \quad (4)$$

where latent variable z is sampled from a fixed latent distribution $p_z(z)$ and real samples x come from a real data distribution $p_{\text{data}}(x)$. The GAN model can be solved by alternating gradient optimisation between the discriminator and the generator. Let $p_G(x)$ be the distribution induced by the generator, for a fixed G , the optimal discriminator is

$$\theta_D^*(x) = \frac{p_{\text{data}}(x)}{p_{\text{data}}(x) + p_G(x)} \quad (5)$$

when the generator distribution exactly matches the data distribution $p_{\text{data}}(x) = p_G(x)$ [52]. To avoid the vanishing gradient problem, when initially the discriminator is very confident and almost always outputs 0, in practice the gradient step for the generator is replaced by

$$\Delta\theta_G = \nabla_{\theta_G} \mathbb{E}_{z \sim p_z(z)} [-\log D_{\theta_D}(G_{\theta_G}(z))]. \quad (6)$$

In so doing, the gradient signals are enhanced, but this is no longer a zero-sum game [53]. When the discriminator is optimal, the minimax game is reduced into a minimisation over the generator only and is equal to

$$\min_{\theta_G} \mathcal{L}(\theta_D^*, \theta_G) = \text{JSD}(p_{\text{data}} || p_G) - \log(2) \quad (7)$$

using the Jensen-Shannon divergence [53]. In practice, optimal θ_D^* is rarely known; thus, minimisation of $\mathcal{L}(\theta_D, \theta_G)$ yields only a lower bound [52], [53].

C. Proposed Method

A GAN can be extended to a conditional model if extra prior information is included to constrain the generator and discriminator; this is known as a conditional GAN [54]. Additional prior information can be discrete labels, text and images [55], [56]. In this study, a GAN conditioned on images was used and Figure 1 shows the overall framework of our conditional GAN-based CS-MRI architecture.

Conditional GAN Loss First, instead of using a CNN, we incorporated the conditional GAN loss into our CS-MRI reconstruction, that is

$$\min_{\theta_G} \max_{\theta_D} \mathcal{L}_{\text{cGAN}}(\theta_D, \theta_G) = \mathbb{E}_{x_t \sim p_{\text{train}}(x_t)} [\log D_{\theta_D}(x_t)] + \mathbb{E}_{x_u \sim p_G(x_u)} [-\log(D_{\theta_D}(G_{\theta_G}(x_u)))], \quad (8)$$

in which there is one input for the generator, i.e., zero-filling reconstruction x_u with aliasing artefacts. After learning, the generator yielded the corresponding de-aliased reconstruction \hat{x}_u , which was fed to the discriminator. The aim is to keep training until the discriminator cannot distinguish a de-aliased reconstruction \hat{x}_u from the fully-sampled ground truth reconstruction x_t . Here x_t and x_u are our input training data or in other words we input x_u conditioned on the given x_t , and output the de-aliased reconstruction \hat{x}_u . Compared to the original conditional GAN [54], in which both the generator and discriminator are conditioned on some extra information, in our DAGAN model only the generator receives the undersampled image input as the conditional information.

DAGAN Architecture Our DAGAN architecture was loosely inspired by [55], [57]. We proposed to use a U-

Net based architecture [47] to construct the generator G that consisted of 8 convolutional layers (encoder layers) and corresponding 8 deconvolutional layers (decoder layers), and each was followed by batch normalisation [58] and leaky ReLU layers. In addition, skip connections were applied to connect mirrored layers between the encoder and decoder paths in order to feed different levels of features to the decoder to gain superior reconstruction details. A hyperbolic tangent function was used as the output activation function for the generator. On the other hand, the discriminator D undertook a classification task to differentiate the de-aliased reconstruction \hat{x}_u from the fully-sampled ground truth reconstruction x_t . It was formed using a standard CNN architecture with 11 convolutional layers, and each was also followed by batch normalisation and leaky ReLU layers. Finally a dense convolutional layer was cascaded and the sigmoid activation function output the classification results (more detail in Supplementary Material).

Content Loss In order to improve the perceptual quality of our reconstruction, a content loss was designed for the training of the generator. This loss consisted of three parts, i.e., a pixel-wise image domain mean square error (MSE) loss, a frequency domain MSE loss and a perceptual VGG loss. First, the MSE based loss functions can be represented as

$$\min_{\theta_G} \mathcal{L}_{\text{iMSE}}(\theta_G) = \frac{1}{2} \|x_t - \hat{x}_u\|_2^2; \quad (9)$$

$$\min_{\theta_G} \mathcal{L}_{\text{fMSE}}(\theta_G) = \frac{1}{2} \|y_t - \hat{y}_u\|_2^2, \quad (10)$$

in which y_t and \hat{y}_u are the corresponding frequency domain data of x_t and \hat{x}_u . The VGG loss is defined as

$$\min_{\theta_G} \mathcal{L}_{\text{VGG}}(\theta_G) = \frac{1}{2} \|f_{\text{vgg}}(x_t) - f_{\text{vgg}}(\hat{x}_u)\|_2^2. \quad (11)$$

Together with the adversarial loss of the generator

$$\min_{\theta_G} \mathcal{L}_{\text{GEN}}(\theta_G) = -\log(D_{\theta_D}(G_{\theta_G}(x_u))), \quad (12)$$

the total loss function can be denoted as

$$\mathcal{L}_{\text{TOTAL}} = \alpha \mathcal{L}_{\text{iMSE}} + \beta \mathcal{L}_{\text{fMSE}} + \gamma \mathcal{L}_{\text{VGG}} + \mathcal{L}_{\text{GEN}}. \quad (13)$$

In this study, normalised MSE (NMSE) was used as the optimisation cost function for the fast CS-MRI reconstruction. However, the solution solely based on the optimisation of the NMSE, which is defined on pixel-wise image difference ($\mathcal{L}_{\text{iMSE}}$), could result in perceptually nonsmooth reconstructions that often lack coherent image details. Therefore, we added NMSE of the frequency domain data as additional constraints ($\mathcal{L}_{\text{fMSE}}$) and also an additional VGG loss (\mathcal{L}_{VGG}) to take the perceptual similarity into account [59]. Once the generator has been trained based on the $\mathcal{L}_{\text{TOTAL}}$, we can apply it to any new inputs (i.e., initial aliased reconstructions after filling zeros into the undersampled k -space), and it will result in the de-aliased reconstruction.

Refinement Learning It is well known that a deep learning based method might be hard to train due to vanishing or exploding gradient problems, however comprehensive parameter tuning may alleviate the problems but subject to large variance of performance with different parameter settings. The original

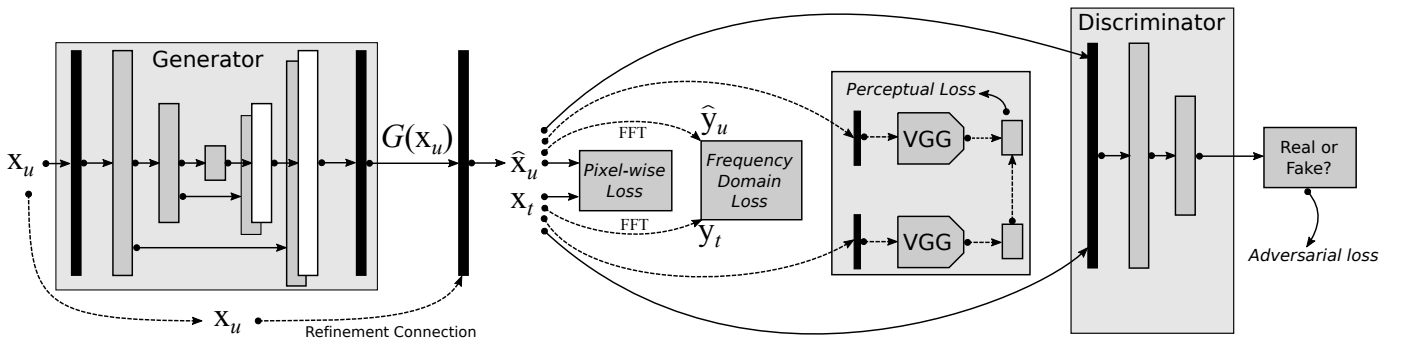


Fig. 1: Schema for our proposed conditional GAN-based de-aliasing for fast CS-MRI (DAGAN).

GAN model is also difficult to train [57] due to the alternating training on the adversarial components. In this study, we proposed a refinement learning to stabilise the training of our DAGAN model, which can yield faster convergence. Essentially, we proposed to use $\hat{x}_u = G_{\theta_G}(x_u) + x_u$ instead of using $\hat{x}_u = G_{\theta_G}(x_u)$. In so doing, we transferred the generator from a conditional generative function to a refinement function, i.e., only generate the missing information, which can dramatically reduce the complexity of the model learning. In addition, in order to ensure that the de-aliased reconstruction \hat{x}_u is in an appropriate intensity scale as the ground truth, we applied a simple ramp function to rescale the image.

Networks and Training Settings The VGG network [49] in this work was pretrained on ImageNet [60]. In particular, we used the conv4 output of the VGG16 as the encoded embedding of the de-aliased output and the ground truth, and computed the MSE between them. We trained separate networks for different undersampling ratios with the following fixed mutual hyperparameters: $\alpha = 15$, $\beta = 0.1$, $\gamma = 0.0025$, initial learning rate of 0.0001, batch size of 25. It is of note that the hyperparameters: α , β and γ are the weights associated with different loss terms. According to previous research [61] and in practice, we found that it is adequate to set these weights such that the magnitude of different loss terms is balanced into similar scales. We adopted Adam optimisation [62] with momentum of 0.5. Each model was learnt by employing early stopping and the learning rate was halved every 5 epochs. Here we want to emphasise that our DAGAN models are robust with minimal parameter tuning so that the same set of hyperparameters were used for our following experiments using different undersampling ratios, various undersampling masks and with and without noises.

Data Augmentation In order to boost the network performance, in addition to conventional data augmentation (e.g., image flipping, rotating, shifting, brightness adjustment and zooming), elastic distortion [63] was also applied to account non-rigid deformation of the imaged organs.

D. Implementation Details and Evaluation Methods

Implementation The implementation of our DAGAN models ¹ has been done using a high-level Python wrapper

(TensorLayer ²) [64] of the TensorFlow ³ library.

Evaluation Methods We report the NMSE, the Peak Signal-to-Noise Ratio (PSNR in dB) and the Structural Similarity Index (SSIM) [59]. The reconstructed fully sampled k -space data was used as ground truth (GT) for validation. In addition to quantitative metrics, we also evaluated our method using qualitative visualisation of the reconstructed MRI images and the error with respect to the GT.

IV. SIMULATION SETTINGS AND RESULTS

A. Simulation Settings

Datasets First, we trained and tested our model using a MICCAI 2013 grand challenge dataset ⁴. We randomly included 100 T₁-weighted MRI datasets for training (70%) and validation (30%) that contained 16095 and 5033 valid 2D images including brain tissues. Independent testing was performed on 50 datasets (contained 9854 2D images). In order to compare with other conventional CS-MRI methods, we randomly chose 50 2D images for comparison. Second, in order to test the diagnostic value of our DAGAN based CS-MRI model, we used the trained model to infer the pathological MRI images, i.e., brain lesion MRI images ⁵ and cardiac MRI images with atrial scarring.

Masks Three different types of undersampling patterns were tested, i.e., 1D Gaussian (**G1D**), 2D Gaussian (**G2D**) and 2D Poisson disc (**P2D**). For the 1D Gaussian mask, 10%, 20%, 30%, 40% and 50% retained raw k -space data were simulated representing 10 \times , 5 \times , 3.3 \times , 2.5 \times and 2 \times accelerations. For 2D Gaussian and 2D Poisson disc masks, 30% retained k -space data was simulated, respectively (For details refer to Supplementary Material).

DAGAN Variations In order to test the effectiveness of different loss components in our cost function, we compared the following DAGAN variations: (1) Pixel-Frequency-Perceptual-GAN-Refinement (**PFPR**): the full model using the GAN architecture with pixel-wise MSE, frequency domain data MSE, VGG loss, and refinement learning; (2) Pixel-Perceptual-GAN-Refinement (**PPGR**): the model without frequency domain data constraints; (3) Pixel-Perceptual-GAN

²<http://tensorlayer.readthedocs.io>

³<https://www.tensorflow.org>

⁴<http://masiweb.vuse.vanderbilt.edu/workshop2013/index.php/>

⁵<http://www.brainlesion-workshop.org/>

¹<https://github.com/nebulaV/DAGAN>

(PPG): the model without refinement learning; (4) Pixel-GAN (PG): the method with pixel-wise MSE only and GAN architecture.

Comparison Methods We compared our DAGAN model with conventional CS-MRI methods including TV [16], SIDWT [65] and RecPF [18], and the state-of-the-art approaches including DLMRI [26], PBDW [23], PANO [66], Noiselet [67], BM3D [50] and DeepADMM [36]. It is worth noticing that all the comparison methods and our DAGAN models were initialised using the baseline zero-filling (ZF) reconstruction to achieve a fair comparison study. The initialisation using a prior reconstructed image (e.g., using SIDWT) may boost the performance of some methods, but can obviously suffer from much longer reconstruction. In addition, we only performed minimum parameter tuning for all the methods. For most of these comparison algorithms (TV, SIDWT, RecPF, PBDW and Noiselet), we used two generic stopping criteria: the number of maximum iterations (800) and the improvement tolerance (0.00001), and the reconstruction stops when either criterion is satisfied. However, the fundamental mechanism of these algorithms are different; therefore, they may have different definitions for the number of iterations. For the DLMRI method, we used the recommended setting, i.e., 40 outer loop iterations with 20 iterations for the K-singular value decomposition algorithm. Similarly, for the BM3D method, we used some recommended number of iterations, e.g., 100. For the PANO method, there is no open source implementation; therefore, we only used the provided executable file to perform the reconstruction (the pre-defined inner loop improvement tolerance is 0.05 [66]). For the DeepADMM method, we performed 500 iterations for the training procedure to avoid possible overfitting and also considered the relatively prolonged training time per iteration. For our DAGAN method we applied the early stopping strategy, which can be considered as an additional and efficient regularisation technique to avoid overfitting [68]. To test the noise tolerance of the CS-MRI methods, we synthesised additive white Gaussian noise, which was added to the k -space before applying the undersampling. Although we assumed that the MRI images before adding noise were clean, the actual data were acquired from MRI scanner that may contain certain amount of noise. In this study, the baseline noise level ($5.5\% \pm 13\%$) was calculated using the method described in [69]. It is of note that although the noise model of magnitude MRI images should follow Rician distribution, the additive Gaussian white noise assumption still holds for the k -space components [70].

B. Results

Comparison of DAGAN Variations Table I tabulates the quantitative comparison results of DAGAN variations (i.e., PG, PPG, PPGR and PFPGR). Overall, the results using 1D Gaussian masks presented in Table I show that adding the refinement learning and frequency domain constraint (PFPGR) improved the average NMSE and PSNR. For all our DAGAN models, we obtained compelling de-aliasing results compared to the ZF reconstruction that contained significant aliasing artefacts. Figure 2 shows example PFPGR reconstructions by

different undersampling ratios (1D Gaussian mask). Qualitatively, there is little difference between the reconstructed images and the GT when the undersampling ratio is $\geq 20\%$. For 10% retained data, most of the aliasing artefacts have still been suppressed effectively and we can still obtain an average PSNR $> 31\text{dB}$, but there is obvious loss of structural details (e.g., organ edges). This is because the k -space is highly undersampled and there is significant information loss in the low frequency regions.

Compared to using pixel-wise MSE only (PG), adding the perceptual loss (PPG) produced similar quantitative results; however, qualitative visualisation showed finer reconstruction details without unrealistic jagged artefacts (Figure 3 (l) vs. (m), (n) and (o)). In addition, a horizontal line profile across a randomly selected case (Figure 4) showed that ZF reconstruction still contained a significant amount of artefacts. PG and PPG clearly reduced the aliasing artefacts, and with refinement learning both PPGR and PFPGR achieve more accurate reconstructed line profile compared to the GT (Figure 4).

Comparison with Other Methods Table I also summarises the comparison results of conventional CS-MRI and some representative state-of-the-art methods. Conventional CS-MRI approaches (TV, SIDWT and RecPF) reconstructed images with limited de-aliasing effect, for example, significant remaining aliasing artefacts can be seen in Figure 3 (c), (d) and (e). Dictionary learning (DLMRI) and patch based methods (PBDW and PANO) obtained better de-aliasing, but with clearly over-smoothed reconstruction details (Figure 3 (f), (g) and (h)). Moreover, there are visible aliasing artefacts in the reconstructed images using the Noiselet method (e.g., in Figure 3 (i)). Although both BM3D and DeepADMM worked quite well (Figure 3 (j) and (k)), all our DAGAN models produced visually more convincing reconstructions with much higher SSIM (Figure 3 (l), (m), (n) and (o)).

Study on Noise and Masks Figures 5 and 6 show the PSNR with respect to different noise levels and various undersampling patterns. Our DAGAN models demonstrated certain tolerance to the noise, e.g., our DAGAN models achieved $> 30\text{dB}$ PSNR when the noise level is $\leq 20\%$. In contrast, other methods were dramatically affected by the noise (Figure 5). For various sampling patterns, our DAGAN models performed better with 2D undersampling masks, and produced superior or comparable reconstruction results with other CS-MRI methods using the same sampling pattern (Figure 6).

Zero-Shot Inference on Pathological Cases Figures 7 and 8 show the reconstruction results using our PFPGR model on example pathological MRI images. It is of note that these results were obtained by using the trained PFPGR model on normal brain MRI images (i.e., randomly selected T₁-weighted MRI data), and there was no pathological MRI image used for training. This can also be referred as a zero-shot inference problem [71]. In general, our DAGAN (PFPGR) model achieved faithful reconstruction with clear pathological patterns been preserved, for example, compared to the ZF reconstruction, our PFPGR model demonstrated superior reconstructed details and better defined brain tumour textures and boundaries (green arrows in Figure 7 (a-c)).

Table I: Quantitative results (NMSE and PSNR) of the comparison study using different random undersampling ratios of the 1D Gaussian mask.

Method	10%		20%		30%		40%		50%		Testing Time (30%)	
	NMSE	PSNR	NMSE	PSNR	NMSE	PSNR	NMSE	PSNR	NMSE	PSNR	CPU (sec) / GPU (ms)	
ZF	0.37±0.09	28.22±4.30	0.21±0.06	34.04±4.58	0.16±0.04	35.61±4.79	0.14±0.04	37.65±4.92	0.10±0.03	40.65±4.95	0.002±0.003	—
TV	0.37±0.09	28.23±4.30	0.20±0.06	34.09±4.57	0.16±0.04	35.68±4.79	0.13±0.04	37.81±4.92	0.09±0.03	40.90±4.95	10.5±1.2	—
SIDWT	0.36±0.10	28.96±4.79	0.15±0.04	36.99±5.34	0.11±0.03	38.93±5.45	0.10±0.03	40.51±6.31	0.07±0.02	43.53±6.26	126.0±22.7	—
RecPF	0.36±0.08	28.62±4.46	0.13±0.02	36.21±5.57	0.11±0.02	37.88±4.90	0.09±0.03	41.37±5.88	0.06±0.02	44.78±6.12	35.9±0.8	—
DLMRI	0.39±0.10	28.25±4.24	0.16±0.03	36.54±4.21	0.12±0.02	38.24±4.29	0.09±0.02	42.11±4.17	0.06±0.01	45.40±4.08	49.1±9.8	—
PBDW	0.21±0.05	31.14±3.65	0.11±0.04	37.56±3.07	0.10±0.04	38.07±3.02	0.07±0.02	41.91±2.95	0.04±0.01	46.05±3.05	223.5±45.5	—
PANO	0.22±0.05	30.84±4.48	0.09±0.02	38.54±4.04	0.08±0.02	39.16±4.02	0.05±0.01	43.05±3.64	0.04±0.01	45.44±3.42	108.6±17.4	—
Noiselet	0.25±0.07	30.19±5.40	0.13±0.04	36.97±5.94	0.10±0.03	39.01±5.96	0.08±0.03	41.51±6.84	0.06±0.02	44.44±6.79	31.1±0.5	—
BM3D	0.20±0.05	31.72±5.23	0.10±0.02	39.40±4.58	0.09±0.02	39.93±4.59	0.07±0.01	43.66±4.48	0.05±0.01	45.89±4.41	87.5±3.0	—
DeepADMM	0.24±0.03	30.70±4.78	0.16±0.09	37.31±3.07	0.15±0.04	37.36±3.27	0.10±0.06	42.36±2.41	0.09±0.04	42.60±2.60	3.2±0.2	—
PG	0.19±0.04	32.82±3.80	0.09±0.02	39.10±3.84	0.09±0.02	39.72±4.30	0.06±0.01	43.28±4.17	0.05±0.01	44.39±4.18	0.2±0.1	5.4±0.1
PPG	0.23±0.06	31.33±2.81	0.10±0.03	38.88±3.41	0.09±0.02	39.78±3.95	0.07±0.03	42.31±2.98	0.05±0.01	45.61±3.73	0.2±0.1	5.4±0.1
PPGR	0.19±0.04	32.84±4.46	0.10±0.02	39.05±3.82	0.09±0.02	39.53±4.12	0.05±0.01	44.45±4.14	0.04±0.01	47.30±3.93	0.2±0.1	5.4±0.1
PFPGR	0.17±0.03	33.79±4.26	0.09±0.02	39.44±3.95	0.08±0.02	40.20±4.07	0.05±0.01	44.83±4.16	0.04±0.01	47.83±4.10	0.2±0.1	5.4±0.1

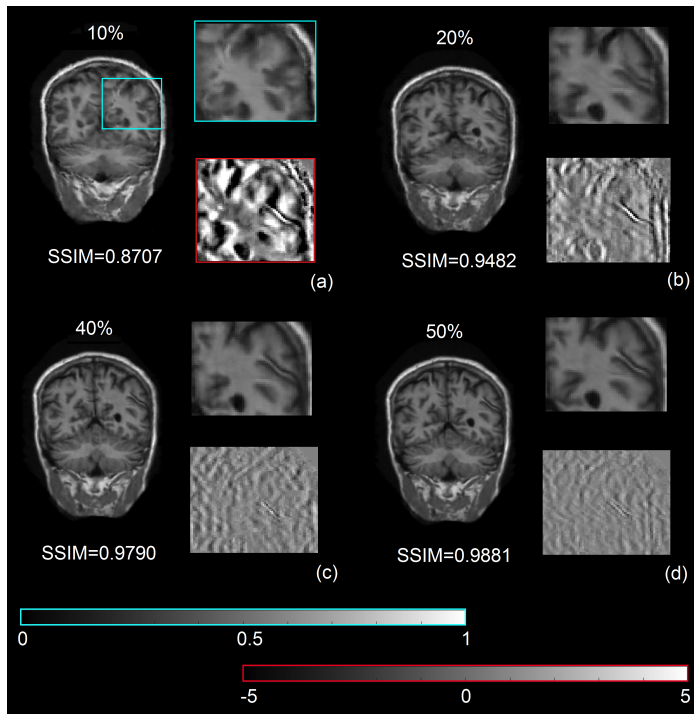


Fig. 2: Qualitative results of using different undersampling ratio. Results were obtained by PFPGR and 1D Gaussian masks with 10% (a), 20% (b), 40% (c) and 50% (d) retained data. GT and results of 30% retained data can be found in Figure 3(a) and (o), respectively. Cyan box in (a): zoomed-in ROI; Red box in (a): zoomed-in difference image ROI (50×) between the reconstructed image and the GT. Corresponding zoomed-in ROIs are shown in (b-d). Colour bars for the image (left) and the difference image (right) are shown at the bottom.

V. DISCUSSION

In this study, we developed a novel conditional GAN based method for fast CS-MRI reconstruction. To the best of our knowledge, the proposed DAGAN model is the first work that incorporates GAN based deep learning for CS-MRI [46]. Overall, our results suggest that the DAGAN model can outperform conventional CS-MRI methods (TV, SIDWT and RecPF) in both qualitative visualisation and quantitative validation. Compared to other state-of-the-art methods (e.g., PANO and BM3D), our DAGAN model can also obtain comparable results. More importantly, the reconstruction time

using DAGAN is much faster than others (about 0.2 sec per 2D image on a CPU or 5 ms on a dedicated GPU) that is feasible for a real-time reconstruction on the MRI scanner.

To emphasise a fair comparison, we initialise all the comparison algorithms with the ZF reconstruction and performed minimum parameter tuning and used generic stopping criteria, e.g., the maximum number of iterations and the improvement tolerance. It should be stressed here that the major purpose of this study is to present our DAGAN model for CS-MRI reconstruction, not benchmarking various reconstruction methods; therefore, a comprehensive comparison of different parameter settings of the compared methods is beyond the scope of the current study. In addition, previous studies have demonstrated that different initialisation could affect the final reconstruction, e.g., the PANO algorithm can perform better with an initialisation using the SIDWT results instead of using ZF [66], but this clearly sacrifices the reconstruction efficiency. Schlemper et al. [43] have also shown that a cascade of alternating CNN and a data consistency layer can achieve superior performance. Such alternating scheme can be also applied to combine our DAGAN model with a data consistency layer or a conventional CS-MRI method; however, this will dramatically reduce the reconstruction efficiency.

Classic fast CS-MRI methods try to solve the image reconstruction using nonlinear optimisation techniques assuming the data is compressible, but normally without considering the prior information of the expected appearance of the anatomy or the possible structure of the undersampling artefacts. This is significantly different from how human radiologists learn to read and interpret MRI images. Radiologists have been trained by reading thousands of MRI images to develop remarkable skills to recognise certain reproducible anatomical and contextual patterns in the images even with known artefacts presented [1], [44]. Our deep learning based DAGAN method aims to imitate this human learning procedure, and therefore shifts the conventional online nonlinear optimisation into an offline training procedure. In other words, our DAGAN method bridges a substantial gap between conventional non-learning methods solving the inverse problem using information from only a single input, and abundant prior knowledge from large training datasets. Compared to our DAGAN method, dictionary learning based methods usually utilise either a fixed

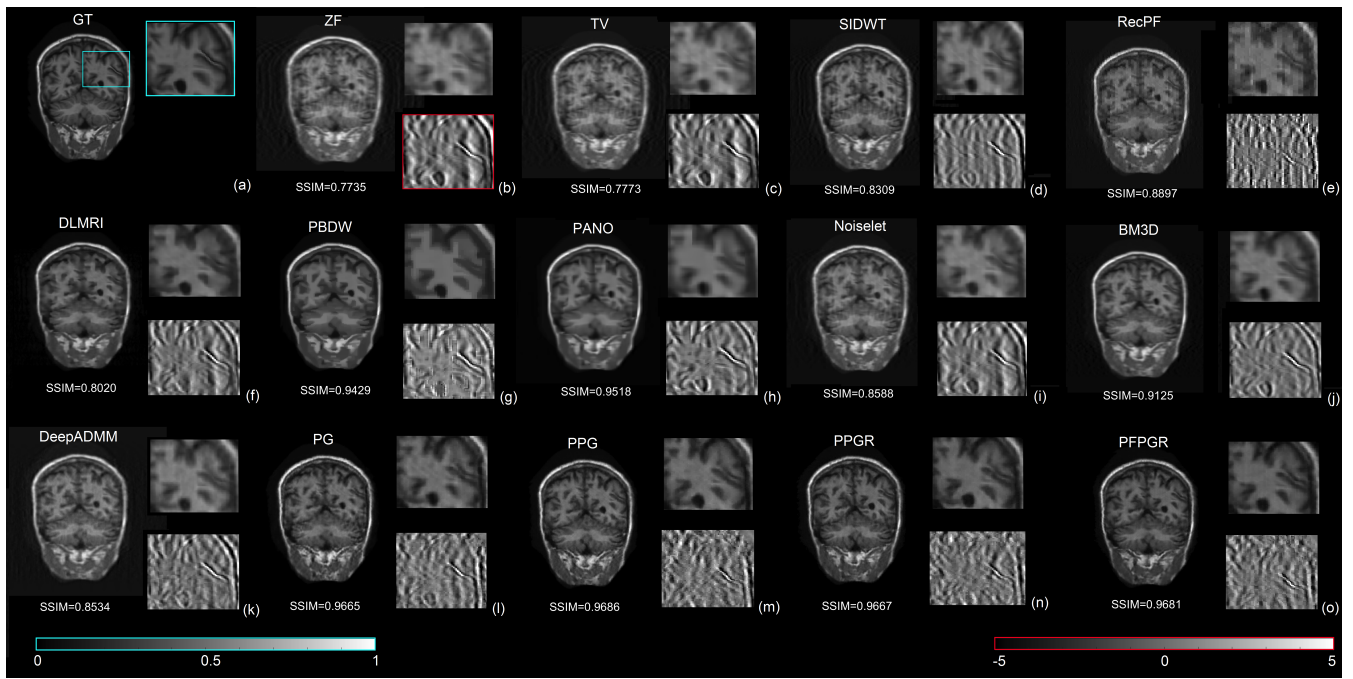


Fig. 3: *Qualitative visualisation.* Cyan boxes in (a): zoomed-in ROI; Red box in (b): zoomed-in difference image ($50\times$) between the reconstructed image and the GT of the same ROI in (a). Corresponding zoomed-in ROIs are shown in (c-o). Results were obtained using 1D Gaussian mask with 30% retained raw data. Colour bars for the image (left) and the difference image (right) are shown at the bottom.




Fig. 4: *Comparison of horizontal line profiles.*

over-complete set of basis or form a dictionary learnt directly from data. For the former learning scheme, there is lack of adaptivity, and for the latter one, the resulting dictionary in sparse coding is not hierarchical as in the deep learning based methods, which in general could provide superior results. In addition, the performance of our DAGAN method is also improved by enriching the training datasets with a comprehensive data augmentation that has not been considered in previous dictionary learning or deep learning based methods [26], [42], [36], [43], [44], [45]. Once a DAGAN model has been trained, it can be used to infer any new input raw data with the




Fig. 5: *PSNR with respect to different noise levels* (1D Gaussian mask was used with 30% retained raw data).

same undersampling ratio. The advantages of the DAGAN model can be two-fold: (1) a more complex nonlinear mapping can be learnt through a comprehensive feature extraction by deep learning, and therefore superior reconstruction details can be obtained; (2) the offline training procedure finishes the labourious optimisation, and the inference avoids any online iterative updating of the reconstruction, which is therefore much more efficient.

In addition to the proposed application of a conditional GAN architecture, a perceptual loss is incorporated to account for the improvement over the reconstructed image quality in terms of the visually more convincing anatomical or pathological details. The idea of the perceptual loss is loosely inspired by GAN based super-resolution [59], which has demonstrated that adding the perceptual loss can achieve better qualitative performance. Our simulation results have also confirmed this,




Fig. 6: PSNR with respect to different undersampling patterns.

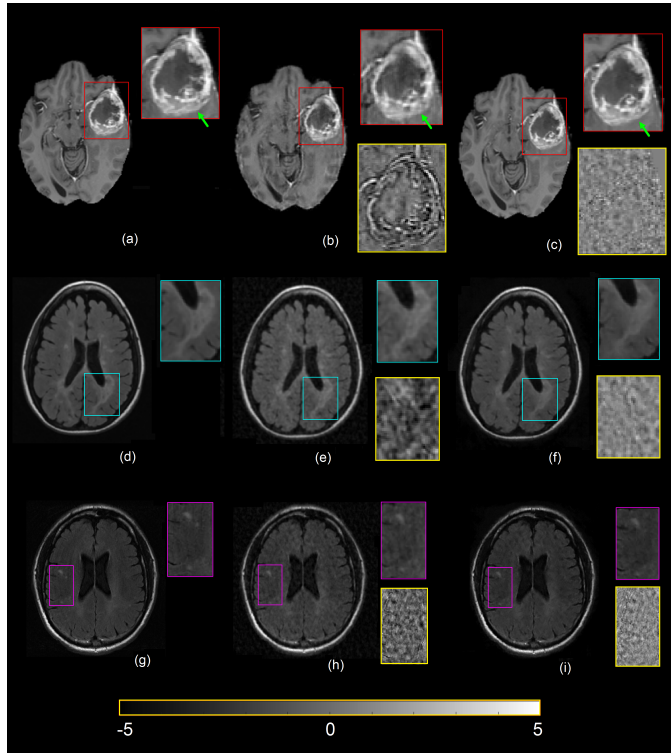


Fig. 7: Zero-Shot inference on pathological MRI images. Left: GT; Middle: ZF; Right: PFPGR. Colour boxes show zoomed-in ROIs. Colour bar below is used for the the difference image (50 \times in yellow box).

which can be attributed to the fact that when reconstructing highly undersampled k -space data, the PG method without perceptual loss can only find an optimal solution to satisfy the MSE criteria but may not perceptually resemble the real data. In addition, our perceptual loss is formed using a pretrained VGG network, which is used to extract high-level features of MRI images; therefore, no *natural-looking* structures have been *hallucinated* in the reconstructed MRI images. Compared to the SR application, CS-MRI solves a more general inverse problem to recover data from undersampled measurements, in which the undersampling pattern is random and noise and artefacts propagation is global due to the frequency domain operation (compared to the regular downsampling pattern and local artefacts in SR). Therefore, the CS-MRI




Fig. 8: Zero-Shot inference on pathological cardiac MRI images. Left: GT; Middle: ZF; Right: PFPGR. Colour boxes show zoomed-in ROIs. Colour bar below is used for the the difference image (50 \times in yellow box).

is a more challenging problem to solve. Furthermore, our DAGAN model can be generalised to solve SR and it is also applicable for solving tomographic reconstruction of other imaging modalities, e.g., Computed Tomography and Positron Emission Tomography.

Together with the perceptual loss, the content loss of our DAGAN model incorporates also MSE loss terms considering both pixel and frequency domain information. In a preliminary study (results shown in Supplementary Material Section 3), the DAGAN model without content loss (using only the adversarial loss) could not achieve acceptable reconstruction. Moreover, the training using only the adversarial loss with refinement learning could not converge. This may be due to the fact that the content loss has provided effective constraints to regularise the generator to synthesise reasonable reconstruction instead of arbitrary image features.

In the context of CS-MRI, one immediate question is whether the conditional GAN architecture would *synthesise* any unrealistic image details in the reconstruction. We studied this by visually scrutinising our reconstruction results, and by a thorough inspection we observed only residual aliasing artefacts when the undersampling ratio is high ($\leq 20\%$); however, there were no unnatural synthesised image details in the intermediate results and final reconstruction (Figure 9). This may be due to the fact that the input of our DAGAN is not totally random, and the ZF reconstruction provided a reasonable initialisation for DAGAN to perform the de-aliasing. In addition, the proposed refinement learning can substantially stabilise the training of the GAN (Figure 10) that is known to be difficult. We noticed that in a previous study on CS based Computed Tomography reconstruction, a similar residual learning method was proposed [72]. It has been used in the U-Net based architecture instead of learning the full-view reconstruction of the Computed Tomography image. Furthermore, their persistent homology analysis demonstrated that the manifold of the full-view reconstruction is topologically more complex than the reconstruction of the residual image, and therefore the residual learning performed better. In our study, a similar refinement learning technique has been applied for training the conditional GAN model. In particular, refinement learning can constrain the generator to reconstruct only

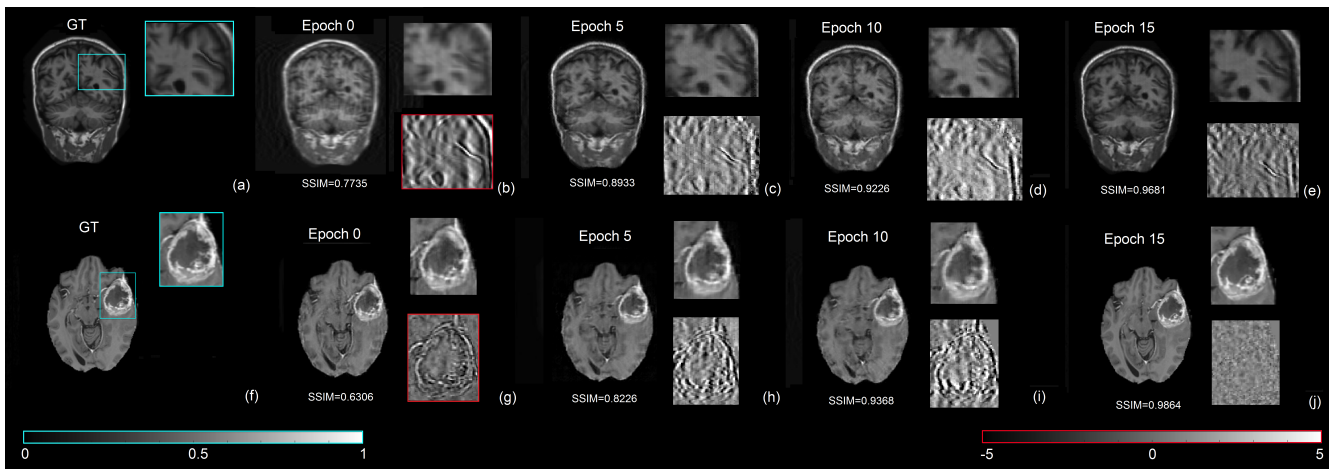


Fig. 9: *Qualitative visualisation of the intermediate results (i.e., epoch 0, 5, 10 and 15). Cyan boxes in (a) a normal brain case and (f) a pathological case: zoomed-in ROI; Red box in (b) and (g): zoomed-in difference image ($50\times$) between the reconstructed image and the GT of the same ROI in (a) and (f). Corresponding zoomed-in ROIs are shown in (c-e) and (h-j). Colour bars for the image (left) and the difference image (right) are shown at the bottom.*

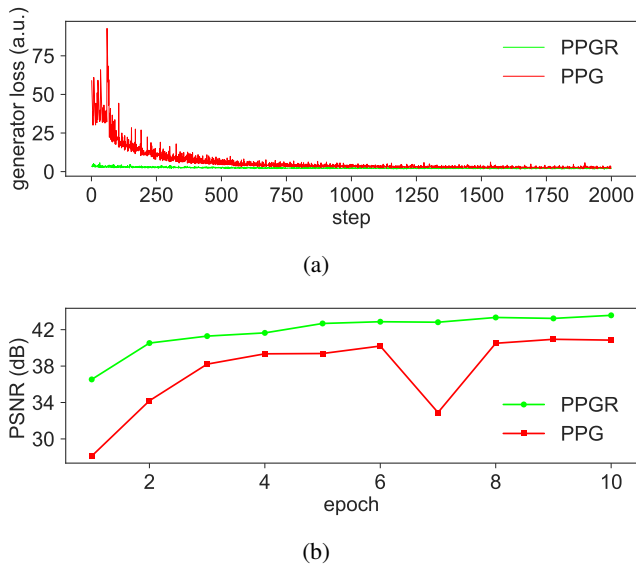


Fig. 10: *Convergence analysis with and without refinement learning.*

the missing details, and prevent it from generating arbitrary features that may not present in real MRI images. In fact, due to limited network capacity and uneven real data distribution, the discriminator might hardly differentiate these unrealistic features, and in turn could incorrectly encourage the generator to reconstruct such arbitrary features. Our convergence analysis (Figure 10) has demonstrated that the proposed refinement learning can dramatically reduce the training complexity of GAN with stabilised and fast-converging training.

In order to demonstrate that the hyperparameters tuning will not affect the DAGAN reconstruction results significantly we performed a robustness analysis by tuning one hyperparameter by $0.5\times$, $5\times$ and $10\times$ of the used setting while fixing the other two. For comparison of different hyperparameter settings, statistical significances were given by a two-sample Wilcoxon rank-sum test. Figure 11 shows that only when we set β to $5\times\beta$, we obtained clearly worse results ($p = 0.012$), but with other parameter settings, there were no significant differences

in the reconstruction results. In addition, our DAGAN method has shown improved tolerance to the additive noise.

Interestingly, the zero-shot inference performed well on pathological cases. The visualisation results showed various brain lesions clearly without any distortion of the lesions or any new lesions being synthesised. It is of note that the training data used for our DAGAN model were normal brain images acquired with a T_1 -weighted MRI sequence while the pathological cases were acquired using different T_1 -weighted or FLAIR sequences. We have shown that our DAGAN model can still reconstruct these images (Figure 7(c), (f), (i)). For cardiac MRI images (Figure 8), the main features of cardiac anatomy were reconstructed reasonably well, although there were some artefacts introduced in the blood pool regions and some loss of fine structural detail. The obvious information loss around the peripheral regions (pink boxes in Figure 8) may be due to the fact that our DAGAN model has been trained using normal brain MRI, and during the inference of cardiac MRI images, the DAGAN model enforces the peripheral regions to be zero as for the brain MRI images. For cardiac MRI, this peripheral information loss is clinically unimportant.

VI. CONCLUSION

The presented study proposes a conditional GAN-based deep learning method for fast CS-MRI reconstruction. The proposed DAGAN method has outperformed conventional CS-MRI approaches and also achieved comparable reconstruction compared to newly developed methods, but the processing time has been remarkably reduced (from seconds to milliseconds per 2D slice) enabling possible real-time application. By combining with existing MRI scanning sequences and parallel imaging, we can envisage this simulation based study to be translated to the real clinical environment.




Fig. 11: PSNR with respect to different parameter settings, and the original parameters are $\alpha = 15$, $\beta = 0.1$, $\gamma = 0.0025$. (1D Gaussian mask was used with 30% retained raw data and reconstructed by PFPGR).

REFERENCES

- [1] K. G. Hollingsworth, "Reducing acquisition time in clinical MRI by data undersampling and compressed sensing reconstruction," *Phys. Med. Biol.*, vol. 60, no. 21, pp. 297–322, 2015.
- [2] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly, "Compressed sensing MRI," *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 72–82, 2008.
- [3] H. Nyquist, "Certain topics in telegraph transmission theory," *Trans. Am. Inst. Elect. Eng.*, vol. 47, no. 2, pp. 617–644, 1928.
- [4] J. Hennig, A. Nauerth, and H. Friedburg, "RARE imaging: a fast imaging method for clinical MR," *Magn. Reson. Med.*, vol. 3, no. 6, pp. 823–833, 1986.
- [5] M. K. Stehling, R. Turner, and P. Mansfield, "Echo-planar imaging: magnetic resonance imaging in a fraction of a second," *Science*, vol. 254, no. 5028, pp. 43–50, 1991.
- [6] G. McGibney, M. R. Smith, S. T. Nichols, and A. Crawley, "Quantitative evaluation of several partial Fourier reconstruction algorithms used in MRI," *Magn. Reson. Med.*, vol. 30, no. 1, pp. 51–59, 1993.
- [7] D. A. Feinberg, J. D. Hale, J. C. Watts, L. Kaufman, and A. Mark, "Halving MR imaging time by conjugation: demonstration at 3.5 kG," *Radiology*, vol. 161, no. 2, pp. 527–531, 1986.
- [8] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, "SENSE: sensitivity encoding for fast MRI," *Magn. Reson. Med.*, vol. 42, no. 5, pp. 952–962, 1999.
- [9] M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase, "Generalized autocalibrating partially parallel acquisitions (GRAPPA)," *Magn. Reson. Med.*, vol. 47, no. 6, pp. 1202–1210, 2002.
- [10] A. Deshmene, V. Gulani, M. A. Griswold, and N. Seiberlich, "Parallel MR imaging," *J. Magn. Reson. Imaging*, vol. 36, no. 1, pp. 55–72, 2012.
- [11] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [12] M. Seeger, H. Nickisch, R. Pohmann, and B. Schölkopf, "Optimization of k-space trajectories for compressed sensing by Bayesian experimental design," *Magn. Reson. Med.*, vol. 63, no. 1, pp. 116–126, 2010.
- [13] H. Chandarana, L. Feng, T. K. Block, A. B. Rosenkrantz, R. P. Lim, J. S. Babb, D. K. Sodickson, and R. Otazo, "Free-breathing contrast-enhanced multiphase MRI of the liver using a combination of compressed sensing, parallel imaging, and golden-angle radial sampling," *Invest. Radiol.*, vol. 48, no. 1, pp. 10–16, 2013.
- [14] E. Levine, B. Daniel, S. Vasanawala, B. Hargreaves, and M. Saranathan, "3D Cartesian MRI with compressed sensing and variable view sharing using complementary poisson-disc sampling," *Magn. Reson. Med.*, vol. 77, no. 5, pp. 1774–1785, 2017.
- [15] M. J. Fair, P. D. Gatehouse, E. V. R. DiBella, and D. N. Firmin, "A review of 3D first-pass, whole-heart, myocardial perfusion cardiovascular magnetic resonance," *J. Cardiovasc. Magn. Reson.*, vol. 17, no. 1, pp. 68, 2015.
- [16] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: the application of compressed sensing for rapid MR imaging," *Magn. Reson. Med.*, vol. 58, no. 6, pp. 1182–1195, 2007.
- [17] K. T. Block, M. Uecker, and J. Frahm, "Undersampled radial MRI with multiple coils. Iterative image reconstruction using a total variation constraint," *Magn. Reson. Med.*, vol. 57, no. 6, pp. 1086–1098, 2007.
- [18] J. Yang, Y. Zhang, and W. Yin, "A fast alternating direction method for TVL1-L2 signal reconstruction from partial fourier data," *IEEE J. Sel. Topics Signal Process.*, vol. 4, no. 2, pp. 288–297, 2010.
- [19] F. Knoll, K. Bredies, T. Pock, and R. Stollberger, "Second order total generalized variation (TGV) for MRI," *Magn. Reson. Med.*, vol. 65, no. 2, pp. 480–491, 2011.
- [20] M. Hong, Y. Yu, H. Wang, F. Liu, and S. Crozier, "Compressed sensing MRI with singular value decomposition-based sparsity basis," *Phys. Med. Biol.*, vol. 56, no. 19, pp. 6311–6325, 2011.
- [21] S. G. Lingala and M. Jacob, "Blind compressive sensing dynamic MRI," *IEEE Trans. Med. Imag.*, vol. 32, no. 6, pp. 1132–1145, 2013.
- [22] Y. Wang and L. Ying, "Undersampled dynamic magnetic resonance imaging using kernel principal component analysis," in *EMBC*, 2014.
- [23] X. Qu, D. Guo, B. Ning, Y. Hou, Y. Lin, S. Cai, and Z. Chen, "Undersampled MRI reconstruction with patch-based directional wavelets," *Magn. Reson. Imag.*, vol. 30, no. 7, pp. 964–977, 2012.
- [24] Z. Zhu, K. Wahid, P. Babyn, and R. Yang, "Compressed sensing-based MRI reconstruction using complex double-density dual-tree DWT," *Int. J. Biomed. Imaging*, vol. 2013, pp. 10, 2013.
- [25] Z. Lai, X. Qu, Y. Liu, D. Guo, J. Ye, Z. Zhan, and Z. Chen, "Image reconstruction of compressed sensing MRI using graph-based redundant wavelet transform," *Med. Image Anal.*, vol. 27, pp. 93–104, 2016.
- [26] S. Ravishanker and Y. Bresler, "MR image reconstruction from highly undersampled k-space data by dictionary learning," *IEEE Trans. Med. Imag.*, vol. 30, no. 5, pp. 1028–1041, 2011.
- [27] J. Caballero, A. N. Price, D. Rueckert, and J. V. Hajnal, "Dictionary learning and time sparsity for dynamic MR data reconstruction," *IEEE Trans. Med. Imag.*, vol. 33, no. 4, pp. 979–994, 2014.
- [28] Z. Zhan, J. Cai, D. Guo, Y. Liu, Z. Chen, and X. Qu, "Fast multiclass dictionaries learning with geometrical directions in MRI reconstruction," *IEEE Trans. Biomed. Eng.*, vol. 63, no. 9, pp. 1850–1861, 2016.
- [29] W. Yin, S. Osher, D. Goldfarb, and J. Darbon, "Bregman iterative algorithms for L1-minimization with applications to compressed sensing," *SIAM J. Imaging Sci.*, vol. 1, no. 1, pp. 143–168, 2008.
- [30] R. Chartrand, "Fast algorithms for nonconvex compressive sensing: MRI reconstruction from very few data," in *ISBI*, 2009.
- [31] O. N. Jaspan, R. Fleisher, and M. L. Lipton, "Compressed sensing MRI: a review of the clinical literature," *Br. J. Radiol.*, vol. 88, no. 1056, pp. 1–12, 2015.
- [32] M. Akçakaya, H. Rayatzadeh, T. Basha, S. N. Hong, R. H. Chan, K. V. Kissing, T. H. Hauser, M. E. Josephson, and W. J. Manning, "Accelerated late gadolinium enhancement cardiac MR imaging with isotropic spatial resolution using compressed sensing: initial experience," *Radiology*, vol. 264, no. 3, pp. 691–699, 2012.
- [33] M. Murphy, M. Alley, J. Demmel, K. Keutzer, S. Vasanawala, and M. Lustig, "Fast L1-SPIRiT compressed sensing parallel imaging MRI: scalable parallel implementation and clinically feasible runtime," *IEEE Trans. Med. Imag.*, vol. 31, no. 6, pp. 1250–1262, 2012.
- [34] X. Qu, W. Zhang, D. Guo, C. Cai, S. Cai, and Z. Chen, "Iterative thresholding compressed sensing MRI based on contourlet transform," *Inverse Probl. Sci. Eng.*, vol. 18, no. 6, pp. 737–758, 2010.

- [35] Y. Liu, J. Cai, Z. Zhan, D. Guo, J. Ye, Z. Chen, and X. Qu, "Balanced sparse model for tight frames in compressed sensing magnetic resonance imaging," *PLOS ONE*, vol. 10, no. 4, pp. e0119584, 2015.
- [36] Y. Yang, J. Sun, H. Li, and Z. Xu, "Deep ADMM-Net for compressive sensing MRI," in *NIPS*, 2016.
- [37] D. Shen, G. Wu, and H. Suk, "Deep learning in medical image analysis," *Annu. Rev. Biomed. Eng.*, vol. 19, pp. 221–248, 2017.
- [38] G. Wang, "A perspective on deep imaging," *IEEE Access*, vol. 4, pp. 8914–8924, 2016.
- [39] C. Dong, C. C. Loy, K. He, and X. Tang, "Image super-resolution using deep convolutional networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 2, pp. 295–307, 2016.
- [40] J. Xie, L. Xu, and E. Chen, "Image denoising and inpainting with deep neural networks," in *NIPS*, 2012.
- [41] F. Agostinelli, M. R. Anderson, and H. Lee, "Adaptive multi-column deep neural networks with application to robust image denoising," in *NIPS*, 2013.
- [42] S. Wang, Z. Su, L. Ying, X. Peng, S. Zhu, F. Liang, D. Feng, and D. Liang, "Accelerating magnetic resonance imaging via deep learning," in *ISBI*, 2016.
- [43] J. Schlemper, J. Caballero, J. V. Hajnal, A. Price, and D. Rueckert, "A deep cascade of convolutional neural networks for dynamic MR image reconstruction," in *IPMI*, 2017.
- [44] K. Hammernik, T. Klatzer, E. Kobler, M. P. Recht, D. K. Sodickson, T. Pock, and F. Knoll, "Learning a variational network for reconstruction of accelerated MRI data," *arXiv preprint arXiv:1704.00447*, 2017.
- [45] D. Lee, J. Yoo, and J. C. Ye, "Deep artifact learning for compressed sensing and parallel MRI," *arXiv preprint arXiv:1703.01120*, 2017.
- [46] S. Yu, et al., "Deep de-aliasing for fast compressive sensing MRI," *arXiv preprint arXiv:1705.07137*, 2017.
- [47] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: convolutional networks for biomedical image segmentation," in *MICCAI*, 2015.
- [48] H. Dong, G. Yang, F. Liu, Y. Mo, and Y. Guo, "Automatic brain tumor detection and segmentation using U-Net based fully convolutional networks," in *MIUA*, 2017.
- [49] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *ICLR*, 2015.
- [50] E. M. Eksioğlu, "Decoupled algorithm for MRI reconstruction using nonlocal block matching model: BM3D-MRI," *J. Math. Imaging Vis.*, vol. 56, no. 3, pp. 430–440, 2016.
- [51] M. J. Ehrhardt and M. M. Betcke, "Multicontrast MRI reconstruction with structure-guided total variation," *SIAM J. Imaging Sci.*, vol. 9, no. 3, pp. 1084–1106, 2016.
- [52] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *NIPS*, 2014.
- [53] M. Arjovsky and L. Bottou, "Towards principled methods for training generative adversarial networks," in *ICLR*, 2017.
- [54] M. Mirza and S. Osindero, "Conditional generative adversarial nets," *arXiv preprint arXiv:1411.1784*, 2014.
- [60] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, and M. Bernstein, "Imagenet large
- [55] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *CVPR*, 2017.
- [56] H. Dong, S. Yu, C. Wu, and Y. Guo, "Semantic image synthesis via adversarial learning," in *ICCV*, 2017.
- [57] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," in *ICLR*, 2016.
- [58] S. Ioffe and C. Szegedy, "Batch normalization: accelerating deep network training by reducing internal covariate shift," in *ICML*, 2015.
- [59] C. Ledig, L. Theis, F. Huszar, J. Caballero, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi, "Photo-realistic single image super-resolution using a generative adversarial network," in *CVPR*, 2017.
- [61] X. Huang, Y. Li, O. Poursaeed, J. Hopcroft, and S. Belongie, "Stacked generative adversarial networks," *arXiv preprint arXiv:1612.04357*, 2016.
- [62] D. Kingma and J. Ba, "Adam: a method for stochastic optimization," in *ICLR*, 2014.
- [63] P. Y. Simard, D. Steinkraus, and J. C. Platt, "Best practices for convolutional neural networks applied to visual document analysis," in *ICDAR*, 2003.
- [64] H. Dong, A. Supratak, L. Mai, F. Liu, A. Oehmichen, S. Yu, and Y. Guo, "TensorLayer: a versatile library for efficient deep learning development," in *ACM-MM*, 2017.
- [65] R. Baraniuk, "Compressive sensing," *IEEE Signal Process. Mag.*, vol. 24, no. 4, pp. 118–121, 2007.
- [66] X. Qu, Y. Hou, F. Lam, D. Guo, J. Zhong, and Z. Chen, "Magnetic resonance image reconstruction from undersampled measurements using a patch-based nonlocal operator," *Med. Image Anal.*, vol. 18, no. 6, pp. 843–856, 2014.
- [67] K. Pawar, G. Egan, and J. Zhang, "Multichannel compressive sensing MRI using noiselet encoding," *PLOS ONE*, vol. 10, no. 5, pp. e0126386, 2015.
- [68] Aurélien Géron, *Hands-on machine learning with Scikit-Learn and TensorFlow: concepts, tools, and techniques to build intelligent systems* No Title, O'Reilly, 2017.
- [69] X. Liu, M. Tanaka, and M. Okutomi, "Single-image noise level estimation for blind denoising," *IEEE Trans. Image Process.*, vol. 22, no. 12, pp. 5226–5237, 2013.
- [70] P. Bao and L. Zhang, "Noise reduction for magnetic resonance images via adaptive multiscale products thresholding," *IEEE Trans. Med. Imag.*, vol. 22, no. 9, pp. 1089–1099, 2003.
- [71] R. Socher, M. Ganjoo, C. D. Manning, and A. Ng, "Zero-shot learning through cross-modal transfer," in *NIPS*, 2013.
- [72] Y. Han, J. Yoo, and J.C. Ye, "Deep Residual Learning for Compressed Sensing CT Reconstruction via Persistent Homology Analysis," *arXiv preprint arXiv:1611.06391*, 2016.