

Recon-GLGAN: A Global-Local context based Generative Adversarial Network for MRI Reconstruction

Balamurali Murugesan ^{*1,2[0000-0002-3002-5845]}, Vijaya Raghavan S², Kaushik Sarveswaran^{2[0000-0002-1525-966X]}, Keerthi Ram², and Mohanasankar Sivaprakasam^{1,2}

¹ Indian Institute of Technology Madras (IITM), India

² Healthcare Technology Innovation Centre (HTIC), IITM, India
 balamurali@htic.iitm.ac.in

Abstract. Magnetic resonance imaging (MRI) is one of the best medical imaging modalities as it offers excellent spatial resolution and soft-tissue contrast. But, the usage of MRI is limited by its slow acquisition time, which makes it expensive and causes patient discomfort. In order to accelerate the acquisition, multiple deep learning networks have been proposed. Recently, Generative Adversarial Networks (GANs) have shown promising results in MRI reconstruction. The drawback with the proposed GAN based methods is it does not incorporate the prior information about the end goal which could help in better reconstruction. For instance, in the case of cardiac MRI, the physician would be interested in the heart region which is of diagnostic relevance while excluding the peripheral regions. In this work, we show that incorporating prior information about a region of interest in the model would offer better performance. Thereby, we propose a novel GAN based architecture, Reconstruction Global-Local GAN (Recon-GLGAN) for MRI reconstruction. The proposed model contains a generator and a context discriminator which incorporates global and local contextual information from images. Our model offers significant performance improvement over the baseline models. Our experiments show that the concept of a context discriminator can be extended to existing GAN based reconstruction models to offer better performance. We also demonstrate that the reconstructions from the proposed method give segmentation results similar to fully sampled images.

Keywords: Magnetic resonance imaging (MRI), Reconstruction · Global Local networks · Segmentation · Deep Learning · Generative adversarial networks · Cardiac MRI

1 Introduction

Medical imaging is the preliminary step in many clinical scenarios. Magnetic resonance imaging (MRI) is one of the leading diagnostic modalities which can

* Code available at <https://github.com/Bala93/Recon-GLGAN>

produce images with excellent spatial resolution and soft tissue contrast. The major advantages of MRI include its non-invasive nature and the fact that it does not use radiation for imaging. However, the major drawback of MRI is the long acquisition time, which causes discomfort to patients and hinders applications in time critical diagnoses. This relatively slow acquisition process could result in significant artefacts due to patient movement and physiological motion. The slow acquisition time of MRI can be attributed 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 by the MRI hardware. In order to accelerate the MRI acquisition process, various methods ranging from Partial Fourier Imaging, Compressed Sensing and Dictionary Learning have been developed [4].

Recently, deep learning based methods have shown superior performance in many computer vision tasks. These methods have been successfully adapted for the MRI reconstruction problem and have shown promising results. The deep learning based methods [9] for MRI reconstruction can be broadly grouped into two : 1) k-space to image domain: the fully sampled image is obtained from zero-filled k-space. Examples include AUTOMAP and ADMM-Net. 2) image to image domain: the fully sampled (FS) image is obtained from the zero-filled (ZF) image. Our focus will be on the models of the latter kind. The work by Wang et al. [14] was the first to use convolutional neural networks to learn the mapping between ZF and FS images. Generative Adversarial Networks (GAN) [6] have shown promising results in many ill-posed inverse problems such as inpainting, super-resolution and denoising when compared to other deep learning based methods. The MRI reconstruction problem, having a similar problem formulation, has been approached with GANs and have shown encouraging results. The main focus of our paper is thus the application of GANs for the MRI reconstruction problem.

In the GANCS work [10], the generator is a residual network, the discriminator is a general deep network classifier and a combination of L1 and adversarial loss constitutes the loss function. Similarly, another work ReconGAN [11] uses a multi-stage network as a generator; a simple deep network classifier for the discriminator, and a combination of MSE loss in the image and frequency domains, adversarial loss constitute the loss function. The addition of the frequency domain loss adds data consistency. DAGAN [15] is another work which uses U-Net as a generator, a deep learning classifier as the discriminator with a combination of MSE loss in the image and frequency domains, adversarial loss and perceptual loss as the loss function. It showed that incorporating the perceptual loss term improved the reconstructed image quality in terms of the visually more convincing anatomical or pathological details. CDFNet [3] proposed the use of a combination of MSE loss in the image and frequency domains along with the Structural Similarity Index Measure (SSIM) as a loss function. This can be extended to a GAN setup. We will refer to this setup as ComGAN. SEGAN [8] proposed a generator network called SU-Net and used a general deep network

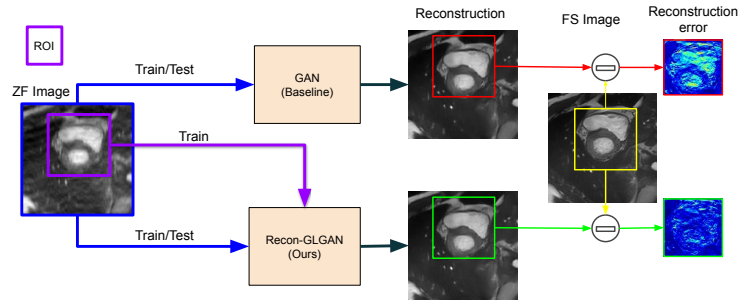


Fig. 1. Illustration depicting the comparison between the baseline GAN model and our Recon-GLGAN model. In the training phase, the ZF image and the ROI are fed in as inputs to the Recon-GLGAN model, while the baseline GAN only takes the ZF image as input. In the testing stage, the ZF image is fed as input to either model to produce the reconstruction (Note: ROI is not used during testing stage). The reconstruction error of the Recon-GLGAN model is lesser than the baseline GAN model in the ROI

classifier as the discriminator. The loss term used is a combination of MSE in the image domain, SSIM and patch correlation regularization.

We refer to the concept of application-driven MRI as described in [2]: incorporating prior information about the end goal in the MRI reconstruction process would likely result in better performance. For instance, in the case of cardiac MRI reconstruction, the physician would be interested in the heart region, which is of diagnostic relevance while excluding the peripheral regions. Using this prior information about the region of interest (ROI) could lead to a better reconstruction. Another perspective is to note that the MRI reconstruction is not the goal in itself, but a means for further processing steps to extract relevant information such as segmentation or tissue characterisation. In general, segmentation algorithms would be interested in the specific ROI. Thus, incorporating prior information about the ROI in the reconstruction process would give two fold benefits : 1) The reconstruction would be better, 2) The segmentation algorithms consequently, could offer better results. The GAN based reconstruction methods described above did not incorporate the application perspective of MRI. Recently, [13] proposed a method in an application-driven MRI context, where the segmentation mask is obtained directly from a ZF image. This work showed encouraging results, but the model produces only the mask as output while the physician would be interested in viewing the FS image. Incorporating the ideas stated above, we propose a novel GAN based approach for MRI reconstruction. A brief outline of our approach compared to baseline GAN approaches is shown in Figure 1. The key contributions of our work can be summarized as follows:

1. We propose a novel GAN architecture, Reconstruction Global-Local GAN (Recon-GLGAN) with a U-Net generator and a context discriminator. The context discriminator consists of a global feature extractor, local feature extractor and a classifier. The context discriminator architecture leverages

global as well as local contextual information from the image. We also propose a loss function which is a linear combination of context adversarial loss and L1 loss in the image domain.

2. We conducted extensive experiments to evaluate the proposed network with a context discriminator for acceleration factors of 2x, 4x and 8x. Our network showed significantly better reconstruction performance when compared with the baseline GAN and UNet architectures for the whole image as well as for a specific region of interest. We also show that the concept of a context discriminator can be easily extended to existing GAN based reconstruction architectures. To this end, we replace the discriminator in the existing GAN based reconstruction architectures with our context discriminator. This showed a significant performance improvement across metrics for an acceleration factor of 4x.
3. We conduct preliminary experiments to show that our model produces reconstructions that result in a better performance for the segmentation task. We demonstrate this using UNet model for segmentation, pre-trained on FS images and the corresponding masks. We observe that the segmentation results produced by the images from our Recon-GLGAN model are similar to FS images in comparison with the ZF and GAN images.

2 Methodology

2.1 Problem Formulation

Let $x_f \in C^N$ be the fully sampled complex image with dimensions $\sqrt{N} \times \sqrt{N}$ arranged in column-wise manner. x_f is obtained from fully sampled k-space measurements ($y_f \in C^N$) through a fully sampled encoding matrix F_f using the relation $y_f = F_f x_f$. During undersampling, a subset of kspace measurements ($y_u \in C^M$) say ($M \ll N$) only are made. This corresponds to an undersampled image x_u by the relation $x_u = F_u^{-1} y_u$. x_u will be aliased due to sub-Nyquist sampling. Reconstructing x_f directly from y_u is ill-posed and direct inversion is not possible due to under-determined nature of system of equations. In our approach, we use deep learning network to learn the mapping between x_u and x_f . The neural network thus learns to minimize the error between predicted fully sampled image (\hat{x}_f) and the ground truth (x_f).

2.2 Generative Adversarial Networks (GAN)

The GAN [6] consists of a generator (G) and discriminator (D). The generator (G) in GAN learns the mapping between two data distributions with the help of discriminator. In the case of MRI reconstruction, the goal of the generator is to learn the mapping between the data distribution of the ZF image (x_u) and FS image (x_f). The discriminator learns to distinguish between the generated and target reconstruction.

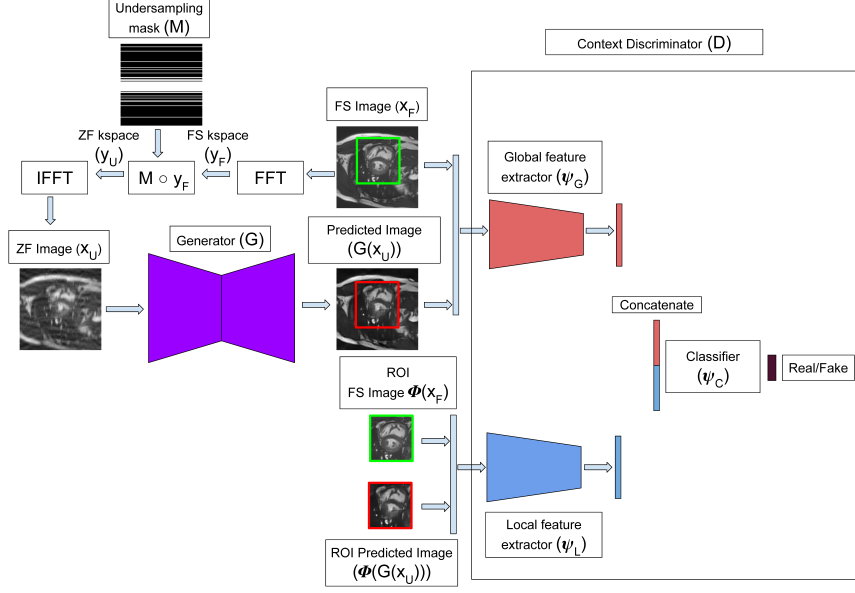


Fig. 2. Recon-GLGAN architecture

2.3 Proposed Reconstruction Global-Local GAN (Recon-GLGAN)

We propose a novel GAN architecture called Reconstruction Global-Local GAN (Recon-GLGAN). The idea is inspired from a GAN based work [5] in the context of image inpainting. The idea behind Recon-GLGAN is to capture both the global and local contextual features. Recon-GLGAN consists of a generator and a context discriminator. The generator (G) tries to learn the mapping between data distribution of ZF image x_u and FS image x_f with the help of the context discriminator which can extract global and local features and classify it as real/fake. The context discriminator consists of three components: global feature extractor, local feature extractor and classifier. The global feature extractor (Ψ_G) takes the entire image as input while the local feature extractor (Ψ_L) takes the region of interest (ROI) (Φ) from the entire image. The classifier network (Ψ_C) takes the concatenated feature vector ($\Psi_G(x) \parallel \Psi_L(x)$) to classify the input image as real/fake. The overview of the proposed architecture is shown in Figure 2. The joint optimization of the generator and context discriminator parameters is given by:

$$\min_{\theta_G} \max_{\theta_D} L_{Recon-GLGAN}(\theta_D, \theta_G) = E_{x_f \sim p_{train}(x_f)} [\log D_{\theta_D}(x_f)] + E_{x_u \sim p_G(x_u)} [-\log(D_{\theta_D}(G_{\theta_G}(x_u)))] \quad (1)$$

$$D_{\theta_D}(x) = \Psi_C(\Psi_G(x) \parallel \Psi_L(\Phi(x))) \quad (2)$$

2.4 Network architecture

Generator (G): The most commonly used encoder-decoder architecture U-Net [12] is used as the generator.

Context Discriminator (D) :

- *Global feature extractor (Ψ_G):* The global feature extractor operates on the whole image. In our case, the input image dimension is 160×160 . The stack of 3 convolutional layers followed by 2 fully connected layers is used as the global feature extractor. Leaky ReLu is used as an activation function for each layer. Average pooling is applied after each convolutional layer. The kernel size of convolutional layer is represented by : (Output channels, Input channels, height, width, stride, padding). The three convolution layers have the following parameters: 1) (32,1,9,9,1,0) 2) (64,32,5,5,1,0) 3) (64,64,5,5,1,0). The 2 fully connected layers converts the feature maps from convolutional layer to 64-dimensional feature vector.
- *Local feature extractor (Ψ_L):* The local feature extractor operates on the specific ROI of an image. In our case, the dimension of the ROI is 60×60 . The architecture is largely similar to that of the global feature extractor except for the dimensions of the feature vector of the fully connected layer, which is modified according to the image dimensions. The output is a 64-dimensional feature vector.
- *Classifier (Ψ_C):* The outputs of the global and the local feature extractors are concatenated together into a single 128-dimensional vector, which is then passed to a single fully-connected layer, to output a single, continuous value. A sigmoid activation function is used so that this value is in the $[0, 1]$ range and represents the probability that the reconstruction is real/fake.

2.5 Loss function

The loss function to accommodate our network design is given below:

$$L_{total} = \lambda_1 L_{imag} + \lambda_2 L_{context} \quad (3)$$

$$L_{imag} = E_{x_u, x_f} [||x_f - G(x_u)||_1] \quad (4)$$

$$L_{context} = E_{x_f} [\log(D(x_f))] + E_{x_u} [-\log(D(G(x_u)))] \quad (5)$$

where L_{imag} is the L1 loss between predicted and target fully sampled image, $L_{context}$ is the context adversarial loss.

3 Experiments and Results

3.1 Dataset

Automated Cardiac Diagnosis Challenge (ACDC) [1] is a cardiac MRI segmentation dataset. The dataset has 150 and 50 patient records for training and

testing respectively. From the patient records, 2D slice images are extracted and cropped to 160×160 . The extracted 2D slices amount to 1841 for training and 1076 for testing. The slices are normalized to the range (0-1). In the context of MRI reconstruction, the slice images are considered as FS images while the ZF images are obtained through cartesian undersampling masks corresponding to 2x, 4x and 8x accelerations.

The MR images in training set have their corresponding segmentation masks whereas the segmentation masks for MR images in test set are not publicly available. The dimensions of the ROI is set to 60×60 based on a study of the sizes of the segmentation masks in the training set. In the training phase, the center of the ROI for each slice is the midpoint of the closest bounding box of the corresponding segmentation mask.

3.2 Evaluation Metrics

Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM), Normalised Mean Square Error (NMSE) metrics are used to evaluate the reconstruction quality for the entire image and its ROI. The segmentation quality is evaluated using Dice similarity coefficient (DICE) and Hausdorff distance (HD).

3.3 Implementation Details

The models were implemented in PyTorch. All models were trained for 150 epochs on two Nvidia GTX-1070 GPUs. Adam optimizer was used for the generator, with a learning rate of 0.001. Stochastic Gradient Descent optimizer was used for the discriminator, with a learning rate of $5e^{-3}$. For the loss term, $\lambda_2 = 4e^{-4}$, and $\lambda_1 = 1$.

The ROI for the MR images in the test set is obtained by following the algorithm described in [7]. This ROI information is not used for inference, it is used only to evaluate the ROI's reconstruction quality.

3.4 Results and Discussion

Reconstruction To evaluate the proposed network, we perform the following experiments:

1) We compare our proposed Recon-GLGAN with the baseline architecture GAN, U-Net, and the ZF images. The metrics for each model for the whole image as well as ROI are shown in Table 1. The results show that our model Recon-GLGAN performs better than the baseline GAN and U-Net across all metrics for all acceleration factors. We also note that our model offers appreciable performance improvement for 4x and 8x acceleration factors compared to 2x. This can be attributed to the fact that the image degradation in the case of 2x is not severe when compared with 4x and 8x. The qualitative comparison of ZF, GAN and Recon-GLGAN for different acceleration factors are shown in Figure 3. In the Figure, it can be observed that reconstruction error of Recon-GLGAN

Table 1. Comparison of Recon-GLGAN with baseline architectures for 2x, 4x and 8x accelerations(FI-Full image)

			NMSE	PSNR	SSIM
2x	FI	Zero-filled	0.01997 ± 0.01	26.59 ± 3.19	0.8332 ± 0.06
		UNet	0.00959 ± 0.00	29.7 ± 2.97	0.9069 ± 0.03
		GAN	0.00958 ± 0.01	29.72 ± 3.03	0.9083 ± 0.03
		Recon-GLGAN	0.00956 ± 0.00	29.74 ± 3.0	0.9108 ± 0.03
	ROI	Zero-filled	0.01949 ± 0.02	25.48 ± 3.73	0.859 ± 0.05
		UNet	0.00952 ± 0.01	28.48 ± 3.03	0.9036 ± 0.04
		GAN	0.00942 ± 0.00	28.53 ± 3.12	0.904 ± 0.04
		Recon-GLGAN	0.00944 ± 0.01	28.54 ± 3.19	0.9065 ± 0.04
4x	FI	Zero-filled	0.03989 ± 0.03	23.65 ± 3.38	0.7327 ± 0.08
		UNet	0.01962 ± 0.01	26.62 ± 3.209	0.8419 ± 0.05
		GAN	0.01934 ± 0.01	26.68 ± 3.08	0.8465 ± 0.05
		Recon-GLGAN	0.01905 ± 0.01	26.8 ± 3.25	0.8497 ± 0.05
	ROI	Zero-filled	0.03886 ± 0.04	22.63 ± 3.87	0.7514 ± 0.07
		UNet	0.01931 ± 0.01	25.46 ± 3.35	0.8242 ± 0.06
		GAN	0.01925 ± 0.02	25.52 ± 3.38	0.8301 ± 0.06
		Recon-GLGAN	0.01878 ± 0.02	25.66 ± 3.26	0.8327 ± 0.06
8x	FI	Zero-filled	0.08296 ± 0.06	20.46 ± 3.24	0.6443 ± 0.09
		UNet	0.03353 ± 0.02	24.26 ± 2.71	0.7547 ± 0.07
		GAN	0.03359 ± 0.02	24.25 ± 2.71	0.7557 ± 0.07
		Recon-GLGAN	0.03286 ± 0.02	24.32 ± 2.68	0.7562 ± 0.07
	ROI	Zero-filled	0.07943 ± 0.08	19.47 ± 3.82	0.6435 ± 0.07
		UNet	0.03147 ± 0.02	23.31 ± 2.88	0.72 ± 0.07
		GAN	0.03129 ± 0.02	23.33 ± 2.92	0.7294 ± 0.07
		Recon-GLGAN	0.03102 ± 0.02	23.34 ± 2.82	0.7293 ± 0.07


**Fig. 3.** From Left to Right: Ground Truth FS image, ZF image, GAN reconstructed image, Recon-GLGAN reconstructed image, ZF reconstruction error, GAN reconstruction error and Recon-GLGAN reconstruction error. From Top to Bottom: Images corresponding to different acceleration factors: 2x, 4x and 8x.

Table 2. GAN based reconstruction architectures and their loss terms

Architecture		Loss function terms
ReconGAN	-	$L_{imag}, L_{global}, L_{freq}$
	GL-ReconGAN	$L_{imag}, L_{context}, L_{freq}$
DAGAN	-	$L_{imag}, L_{global}, L_{freq}, L_{vgg}$
	GL-DAGAN	$L_{imag}, L_{context}, L_{freq}, L_{vgg}$
SEGAN	-	$L_{imag}, L_{global}, L_{ssim}$
	GL-SEGAN	$L_{imag}, L_{context}, L_{ssim}$
COMGAN	-	$L_{imag}, L_{freq}, L_{global}, L_{ssim}$
	GL-COMGAN	$L_{imag}, L_{freq}, L_{context}, L_{ssim}$

for entire image and its ROI is better than GAN. But, it is evident that, the reconstruction error of Recon-GLGAN is significantly better than GAN in the ROI compared with the entire image. This behaviour can be attributed to the design of context discriminator which has a separate feature extraction path for specified ROI. The design of context discriminator enables the generator to specifically learn the ROI along with the entire image during the training phase.

Table 3. Reconstruction metric comparison for full image and region of interest for various GAN based reconstruction architecture for 4x accelerations(FI - Full Image)

			NMSE	PSNR	SSIM
ReconGAN	FI	-	0.01857 ± 0.01	26.82 ± 2.89	0.8485 ± 0.05
		GL-ReconGAN	0.01844 ± 0.01	26.91 ± 3.12	0.8498 ± 0.05
	ROI	-	0.018 ± 0.01	25.76 ± 3.06	0.832 ± 0.06
		GL-ReconGAN	0.01836 ± 0.01	25.72 ± 3.24	0.8336 ± 0.06
SEGAN	FI	-	0.01862 ± 0.01	26.84 ± 3.10	0.8483 ± 0.06
		GL-SEGAN	0.01817 ± 0.01	27.02 ± 3.4	0.8545 ± 0.05
	ROI	-	0.0185 ± 0.01	25.64 ± 3.19	0.8308 ± 0.07
		GL-SEGAN	0.01793 ± 0.01	25.87 ± 3.56	0.838 ± 0.06
ComGAN	FI	-	0.01899 ± 0.01	26.78 ± 3.14	0.8481 ± 0.05
		GL-ComGAN	0.01789 ± 0.01	27.06 ± 3.26	0.8505 ± 0.05
	ROI	-	0.01872 ± 0.01	25.64 ± 3.28	0.8315 ± 0.06
		GL-ComGAN	0.01766 ± 0.02	25.91 ± 3.25	0.834 ± 0.06
DAGAN	FI	-	0.01903 ± 0.01	26.75 ± 3.06	0.8452 ± 0.06
		GL-DAGAN	0.01851 ± 0.01	26.87 ± 3.03	0.845 ± 0.06
	ROI	-	0.01838 ± 0.01	25.68 ± 3.04	0.8272 ± 0.07
		GL-DAGAN	0.01858 ± 0.01	25.62 ± 3.016	0.8277 ± 0.07

2) We attempt to show that the concept of a context discriminator can be extended to existing GAN based works for MRI reconstruction. The different GAN based architectures and their corresponding loss functions can be found in Table 2. In this experiment to ensure a fair comparison, the generator is set to U-Net, discriminator is set to global feature extractor(Ψ_G) followed by a




Fig. 4. From Left to Right: Ground Truth FS image, ZF image for 4x undersampling factor, GAN with basic discriminator reconstructed image, GAN with context discriminator reconstructed image, ZF reconstruction error, GAN with basic discriminator reconstruction error and GAN with context discriminator reconstruction error. From top to bottom: ReconGAN, SEGAN, ComGAN, DAGAN.

classifier(Ψ_C)(basic discriminator) and the loss functions are taken from their respective works [15,3,8,10]. This arrangement means that the difference between the various GAN based architectures comes only from the generator loss. In this experiment, we replace the basic discriminator of the GAN architectures with our proposed context discriminator. The results comparing the GAN architectures with basic discriminator and context discriminator are reported in Table 3. From the Table, it is clear that the GAN with context discriminator have shown improved results compared to GAN with basic discriminator for different generator loss. A few sample results comparing the GAN based reconstruction methods with basic and context discriminator are shown in Figure 4. From the figure we observe that the ROI's reconstruction error for GAN with context discriminator is lesser compared to GAN with the basic discriminator. This shows that the context discriminator can be extended to other GAN based reconstruction methods.

Segmentation Image segmentation is an important task in medical imaging and diagnosis. For instance, in the case of cardiac MRI, the segmentation of left ventricle (LV), right ventricle (RV) and myocardium (MC) are used for cardiac function analysis. Advances in deep learning networks have produced state-of-the-art results. These networks are trained on the FS images and, testing the network with ZF images will result in an unsatisfactory segmentation. We note that a better reconstruction, which is close to the FS image would result in better segmentation performance. In this experiment, we would like to show that the segmentation performance on the reconstructed images from our Recon-GLGAN model is better than the baseline GAN model. To demonstrate this, we use the




Fig. 5. From Left to Right: FS image, ZF image, GAN reconstructed image, Recon-GLGAN reconstructed image, Ground Truth FS segmentation mask, Segmentation mask for ZF, Segmentation mask for GAN reconstructed image and Segmentation mask for Recon-GLGAN reconstructed image. From top to bottom: Sample 1, 2 and 3

most widely used segmentation network U-Net [12]. U-Net is trained on the FS images to produce multi-class (LV, RV and MC) segmentation outputs. Since the ground truth segmentation masks are unavailable for the test set of the ACDC dataset, we instead use the outputs of the FS images in the test set as ground truth. The reconstructed images from GAN and Recon-GLGAN are passed to the UNet and the corresponding segmentation masks are obtained. The obtained segmentation masks for sample images are shown in Figure 5. It is evident from the figure that our network’s performance is closest to FS followed by GAN and ZF images. The same are quantified using the segmentation metrics Dice and Hausdorff for the sample images in Figure 6.




Fig. 6. Segmentation metrics: Dice and HD comparison for image samples 1, 2 and 3

4 Conclusion

In this work, we proposed a novel GAN network, Recon-GLGAN. The context discriminator proposed in Recon-GLGAN helps to capture both global and local features enabling a better overall reconstruction. We showed the extensibility of our discriminator with various GAN based reconstruction networks. We also demonstrated that the images obtained from our method gave segmentation results close to fully sampled images.

References

1. Bernard, O., Lalande, A., Zotti, C., Cervenansky, F., et al.: Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: Is the problem solved? *IEEE Transactions on Medical Imaging* **37**(11), 2514–2525 (Nov 2018)
2. Caballero, J., Bai, W., Price, A.N., Rueckert, D., Hajnal, J.V.: Application-driven mri: Joint reconstruction and segmentation from undersampled mri data. In: Golland, P., Hata, N., Barillot, C., Hornegger, J., Howe, R. (eds.) *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2014*. pp. 106–113 (2014)
3. Dedmari, M.A., Conjeti, S., Estrada, S., Ehses, P., Stöcker, T., Reuter, M.: Complex fully convolutional neural networks for mr image reconstruction. In: *Machine Learning for Medical Image Reconstruction*. pp. 30–38 (2018)
4. Hollingsworth, K.G.: Reducing acquisition time in clinical MRI by data under-sampling and compressed sensing reconstruction. *Physics in Medicine and Biology* **60**(21), R297–R322 (oct 2015)
5. Iizuka, S., Simo-Serra, E., Ishikawa, H.: Globally and locally consistent image completion. *ACM Trans. Graph.* **36**(4), 107:1–107:14 (Jul 2017)
6. Isola, P., Zhu, J., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 5967–5976 (July 2017)
7. Khened, M., Kollerathu, V.A., Krishnamurthi, G.: Fully convolutional multi-scale residual densenets for cardiac segmentation and automated cardiac diagnosis using ensemble of classifiers. *Medical Image Analysis* **51**, 21 – 45 (2019)
8. Li, Z., Zhang, T., Zhang, D.: SEGAN: structure-enhanced generative adversarial network for compressed sensing MRI reconstruction. *CoRR* **abs/1902.06455** (2019)
9. Lundervold, A.S., Lundervold, A.: An overview of deep learning in medical imaging focusing on mri. *Zeitschrift fr Medizinische Physik* **29**(2), 102 – 127 (2019)
10. Mardani, M., Gong, E., Cheng, J.Y., Vasanawala, S.S., Zaharchuk, G., Xing, L., Pauly, J.M.: Deep generative adversarial neural networks for compressive sensing mri. *IEEE Transactions on Medical Imaging* **38**(1), 167–179 (Jan 2019)
11. Quan, T.M., Nguyen-Duc, T., Jeong, W.: Compressed sensing mri reconstruction using a generative adversarial network with a cyclic loss. *IEEE Transactions on Medical Imaging* **37**(6), 1488–1497 (June 2018)
12. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional Networks for Biomedical Image Segmentation. In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*. pp. 234–241 (2015)
13. Schlemper, J., Oktay, O., Bai, W., Castro, D.C., Duan, J., Qin, C., Hajnal, J.V., Rueckert, D.: Cardiac mr segmentation from undersampled k-space using deep latent representation learning. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*. pp. 259–267 (2018)
14. Wang, S., Su, Z., Ying, L., Peng, X., Zhu, S., Liang, F., Feng, D., Liang, D.: Accelerating magnetic resonance imaging via deep learning. In: *2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI)*. pp. 514–517 (April 2016)
15. Yang, G., Yu, S., Dong, H., Slabaugh, G., Dragotti, P.L., Ye, X., Liu, F., Arridge, S., Keegan, J., Guo, Y., Firmin, D.: Dagan: Deep de-aliasing generative adversarial networks for fast compressed sensing mri reconstruction. *IEEE Transactions on Medical Imaging* **37**(6), 1310–1321 (June 2018)