

# Cine cardiac MRI reconstruction using a convolutional recurrent network with refinement

Yuyang Xue<sup>1\*</sup>, Yuning Du<sup>1\*</sup>, Gianluca Carloni<sup>2,3\*</sup>, Eva Pachetti<sup>2,3\*</sup>, Connor Jordan<sup>4\*</sup>, and Sotirios A. Tsaftaris<sup>1,5</sup>

<sup>1</sup> Institute for Digital Communications, University of Edinburgh, United Kingdom

<sup>2</sup> Institute of Information Science and Technologies “Alessandro Faedo”, National Research Council of Italy (CNR), Pisa, Italy

<sup>3</sup> Department of Information Engineering, University of Pisa, Italy

<sup>4</sup> Institute for Infrastructure and Environment, University of Edinburgh, United Kingdom

<sup>5</sup> The Alan Turing Institute, London, United Kingdom

**Abstract.** Cine Magnetic Resonance Imaging (MRI) allows for understanding of the heart’s function and condition in a non-invasive manner. Undersampling of the  $k$ -space is employed to reduce the scan duration, thus increasing patient comfort and reducing the risk of motion artefacts, at the cost of reduced image quality. In this challenge paper, we investigate the use of a convolutional recurrent neural network (CRNN) architecture to exploit temporal correlations in supervised cine cardiac MRI reconstruction. This is combined with a single-image super-resolution refinement module to improve single coil reconstruction by 4.4% in structural similarity and 3.9% in normalised mean square error compared to a plain CRNN implementation. We deploy a high-pass filter to our  $\ell_1$  loss to allow greater emphasis on high-frequency details which are missing in the original data. The proposed model demonstrates considerable enhancements compared to the baseline case and holds promising potential for further improving cardiac MRI reconstruction.

**Keywords:** Cardiac MRI Reconstruction · MRI Acceleration · MRI Refinement · CRNN

## 1 Introduction

Cardiac magnetic resonance imaging is a powerful, non-invasive tool to aid visualise the heart’s chambers, valves, blood vessels and surrounding tissue. To gain a 3D depiction of the heart, a sequential acquisition process of 2D slices is used, with the scanning duration increasing with the number of slices and temporal resolution desired. Thus for detailed scanning, multiple cardiac cycles must be monitored and the duration of the MRI process can consequently exceed the patients ability to remain steady and hold their breath. By undersampling in the  $k$ -space data acquisition process, the scan time can be substantially reduced

---

\* These authors contributed equally.

at the cost of missing information that must be interpolated. Deep learning achieves  $k$ -space reconstruction with greater prior knowledge for the regularisation term that covers the missing  $k$ -space domain, without requiring an iterative optimisation process and hence greatly accelerating the reconstruction rate.

Various architectures have been explored for MRI reconstruction, including convolutional neural networks (CNNs) and U-Nets [7,8,12,21], variational networks [5] and generative adversarial networks [13,22]. Other deep learning methods that exploit prior knowledge and extend traditional iterative methods include the model-based deep-learning architecture [1] and deep density priors [18]. Enhancing cine MRI through deep learning involves not only capitalising on the spatial relationships acquired from a given dataset but also leveraging temporal correlations. This has been evidenced across various model architectures [10,13,19,25] as well as through registration-based [23] and motion-guided alignment [6] approaches. In [16], data sharing layers were incorporated in a cascaded CNN, whereby adjacent time step  $k$ -space data was used to fill the unsampled lines. In [15], recurrent connections are employed across each iteration step as well as bidirectional convolutional recurrent units facilitating knowledge sharing between iterations and input time frames, respectively.

Working within the confines of the challenge, we explored various architectures and found the CRNN block of [15] to perform best within the given limitations in memory and reconstruction time set by the organisers. This was subsequently combined with a lightweight refinement module inspired by single-image super-resolution approaches [3] to perform further de-noising and resolve finer details. The rest of the paper is organised as follows: Sections 2 and 3 describes the problem, dataset, and methodology, Sect. 4 presents the results of experiments and Sect. 5 and 6 provide discussion and conclusions, respectively.

## 2 Problem formulation and dataset

The objective of MRI reconstruction is to address an ill-posed inverse problem, retrieving image information denoted as  $\mathbf{x} \in \mathbb{C}^N$  from acquired undersampled signals  $\mathbf{y} \in \mathbb{C}^K$ , where  $K \ll N$ . This procedure can be depicted using a linear forward operator  $\mathbf{E}$ , which defines the characteristics of the forward problem:

$$\mathbf{y} = \mathbf{E}\mathbf{x} + \epsilon. \quad (1)$$

Eq. 1 represents the general form of MRI reconstruction. The goal of reconstruction is to minimise the difference between  $\mathbf{x}$  and the ground truth. Therefore, the reconstruction problem can be defined as follows:

$$\tilde{\mathbf{x}} = \arg \min_{\mathbf{x} \in \mathbb{C}^N} \frac{\lambda}{2} \|\mathbf{E}\mathbf{x} - \mathbf{y}\| + f_{\theta}(\mathbf{x}). \quad (2)$$

Here,  $f_{\theta}$  denotes a neural network for image reconstruction with trainable parameters  $\theta$  and  $\lambda$  controls the balance between the network and data consistency.

**Data** Our model is evaluated on the CMRxRecon Challenge dataset from the 26<sup>th</sup> International Conference on Medical Image Computing and Computer Assisted Intervention. The dataset includes both short-axis (SA) and long-axis (LA) (two-chamber, three-chamber and four-chamber) views under acceleration rates of  $4\times$ ,  $8\times$ ,  $10\times$ . The dataset was obtained following recommended protocols and processing [20,24], more details of which can be found on the challenge website [4]. The 300-patient dataset is split 120:60:120 between challenge training, validation, and testing respectively. Only the challenge training set contained ground truth reference data, hence this was further split 90:20:10 for training, evaluation, and testing respectively for all models.

**Data pre-processing** The unpadded image size varies between widths of 132, 162, 204 & 246 and heights of 448 & 512 pixels. To maintain consistent input size, we apply zero-padding for image sizes of  $256 \times 512$  following the Inverse Fourier Fast Transform, with the outputs cropped after inference.

When using the approach from [10], the computationally intensive conjugate gradient step was a limiting factor due to the GPU’s initial 24GB memory constraint set by the challenge organisers. We thus chose to use the coil combined data rather than the multi-coil format, which allowed use of a simpler data consistency step, at a potential loss of accuracy without using the extra information. Likewise, we used a single channel for the processed image instead of using independent channels for amplitude/phase or real/imaginary components, as adopted by [5,7,12,16].

The SA data are 3-dimensional spatially with an additional time component. It is therefore conceivable that full 4D convolutional kernels could be used to fully utilise spatio-temporal redundancies, but this would result in extremely large memory requirements as discussed in [11]. Furthermore, studies such as [19] have demonstrated that it is preferable to have a larger  $2D + t$  network than a smaller 3D-input network with equivalent memory consumption. Therefore, due to the large image size, we choose to use time-series batches of 2D depth slices as per [15] rather than  $3D + t$  or 4D for the long-axis images.

### 3 Methodology

#### 3.1 Model exploration

The initial limitations for inference imposed by the organisers were 24GB GPU VRAM and 4 hours for the reconstruction of the test dataset, which was later increased after our initial investigations. Pre-trained models or loss functions were not permitted. Denoising diffusion probabilistic models (DDPM) were found to take too long in inference, whilst transformer models have been found to lead to heavily pixelated reconstructions. Hence, more conventional approaches were tested, building upon an existing repository<sup>6</sup>.

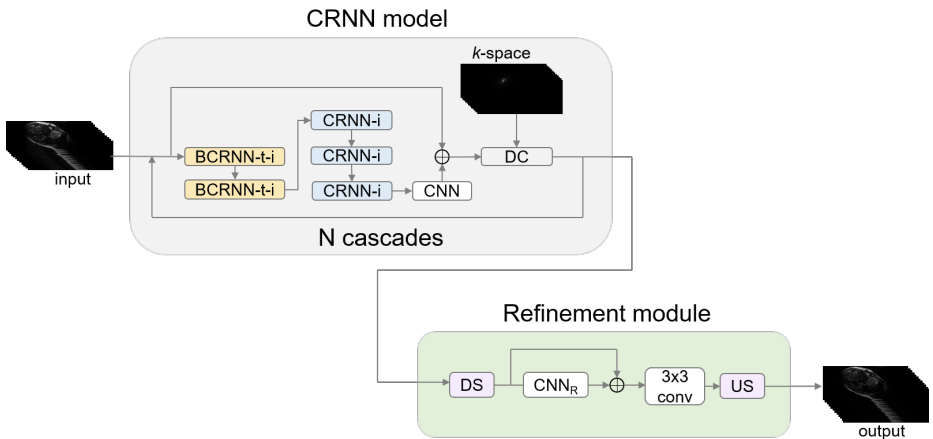
<sup>6</sup> <https://github.com/f78bono/deep-cine-cardiac-mri>

We compared networks similar to the CineNet [10] and CRNN networks [15]. The number of parameters in each network were maximised such that the full 24GB VRAM memory would be used in training. A 2D U-Net is deployed to serve as an additional baseline to compare all models to. The U-Net is trained on a slice-by-slice basis with 3 cascades and 48 feature map channels. Weight sharing is used when training the model, and  $\lambda$  is set to be learnable with an initialisation of  $\log(10^{-1})$ . The learning rate is  $3 \times 10^{-4}$ , and the Adam optimiser is deployed to guide the training process.

### 3.2 Model architecture

A high-level depiction of the complete architecture of the final model is presented in Fig. 1. The backbone of proposed architecture is based on the CRNN block detailed in [15]. The first step in CRNN is a bidirectional convolutional recurrent unit (BCRNN) with three convolution layers: a standard convolution between layers, one convolution between temporal slices, and one between iterations. This is followed by three convolutional recurrent units (CRNN) which evolve only over iterations before a plain CNN. Finally, residual connections are employed prior to a data consistency term, preserving the information from sampled data.

Aiming to improve performance, we include an additional BCRNN unit to further exploit spatio-temporal dependencies, followed by a refinement module to further denoise the output of the CRNN model and refine further details. We deploy a very lightweight single-image super-resolution network, Bicubic++ [3], which maintains short reconstruction times. The refinement module first learns lower resolution features to decrease computational cost and then performs numerous convolutions to denoise the image before a final convolutional filter



**Fig. 1.** Final model architecture: BCRNN, CRNN, and CNN units with a data consistency (DC) step from [15] for primary reconstruction. "t" and "i" denote time and iterations, respectively. The low-cost refinement module, inspired by [3], includes downsampling (DS), CNN, and upsampling (US) units.

and upscaling back to the original image size. We test the performance of the network with end-to-end and separate learning for each module.

**Loss function** In the context of image reconstruction,  $\ell_1$  loss,  $\ell_2$  loss and SSIM loss are widely used to constrain models for high-quality reconstructed images, but often disregard the complex nature of MRI data. Thus, we investigate a range of losses using an additional loss term, denoted the  $\perp$ -loss [17]. The  $\perp$ -loss adds a phase term which can be combined with  $\ell_1$ ,  $\ell_2$  or SSIM losses to address the asymmetry in the magnitude/phase loss landscape. This operates on the polar representation of complex numbers, rather than on two real value channels for magnitude and phase, thus taking advantage of the fact that fully symmetric loss functions can improve task performance [14]. For the separate training of the CRNN and the refinement module,  $\perp$ -losses are only utilised for the CRNN output, with  $\ell_1$  and SSIM loss functions deployed for the refinement module. For the end-to-end training,  $\ell_1$  and SSIM loss are employed to constrain both CRNN and refinement module. We split the  $\ell_1$  loss by introducing a high-pass frequency filter, allowing us to emphasise the high-frequency content in our reconstructed images to resolve finer details. We denote this as  $\ell_{1\text{split}}$ .

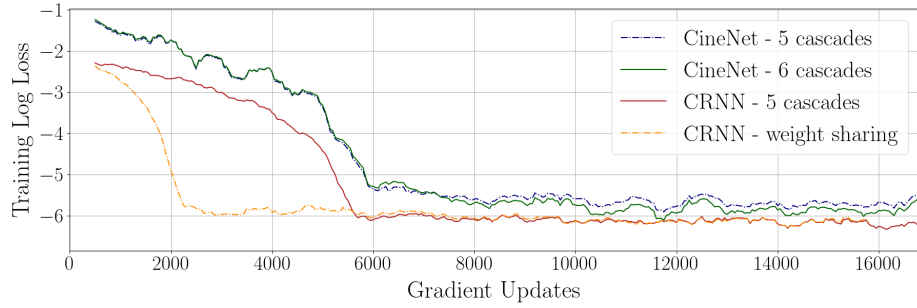
In training, losses were quantified across the entire image, whereas for the validation leaderboard, assessment was limited to the initial 3 time frames and the central sixth portion of the images. The competition metrics were structural similarity index measure (SSIM), normalised mean square error (NMSE) and peak signal-to-noise ratio (PSNR). Hence, whilst the complete reconstructed images often surpassed SSIM values of 0.98, validation scores only reached 0.85.

### 3.3 Implementation details

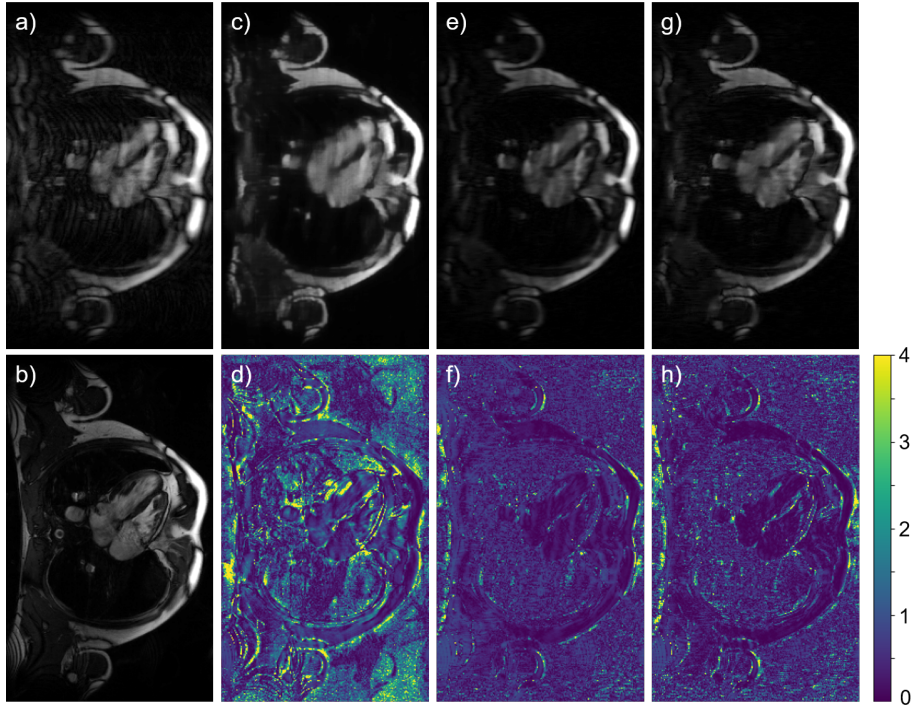
Implementation details and our code is available at: [https://github.com/vios-s/CMRxRECON\\_Challenge\\_EDIPO](https://github.com/vios-s/CMRxRECON_Challenge_EDIPO)

## 4 Results

**Model choice and weight sharing** Fig. 2 shows the training losses between models, demonstrating the stronger performance of the CRNN model. The point of convergence for the CRNN with weight-sharing and the CineNet models is similar, but the CRNN networks start at a much lower loss value. This is despite the non-weight-sharing model (1.1M) having over  $2\times$  more trainable parameters than the 6-cascade CineNet model (0.5M). Between CRNN models, we see more rapid convergence in the weight-sharing model as there are less parameters to optimise and reduced likelihood of early overfitting. However, the lower number of parameters leads to reduced expressive power and is outperformed by the non-weight-sharing model with insufficient memory gains to justify its use. Using the  $\perp$ -loss only, the weight-sharing model had SSIM of 0.683, NMSE of 0.123 and PSNR of 23.917, performing notably worse than the non-weight-sharing model, as presented in the next section. Fig. 3 shows the reconstruction through the CineNet and the CRNN (with and without weight-sharing) models.



**Fig. 2.** Log loss during exploratory training of modified CineNet and CRNN (with and without weight-sharing between kernels). Note that the implementation is not identical to the original works.



**Fig. 3.** Reconstruction (top) and associated error maps (bottom) for the initial network investigation. (a)  $8 \times$  undersampled LAX input (b) fully sampled ground truth (c,d) CineNet model (6 cascades) (e,f) CRNN model (weight-sharing between cascades) (g,h) CRNN model (no weight-sharing).

**Loss function investigation** Table 1 presents the findings of investigations of the loss functions using a low-cost CRNN model. Use of the  $\perp$ - $\ell_1$  loss led to an improvement in SSIM and PSNR compared to  $\ell_1$  loss alone, but a slightly higher NMSE error. Providing greater emphasis on high-frequency data using a high-pass filter led to improved SSIM but lower NMSE and PSNR performance.

**Table 1.** Performance comparisons for different loss function combinations, evaluated on the validation data. A 48-channel 5 cascade CRNN network was used without the refinement module.  $\ell_{1\text{split}}$  denotes the  $\ell_1$  loss whereby a high-pass filter is used to provide more focus on the high frequency content, in addition to the conventional  $\ell_1$  loss.

Metric	Loss function				
	$\perp$	$\ell_1$	$\perp\text{-}\ell_1$	$\perp\text{-}\ell_{1\text{split}}$	$\perp\text{-SSIM-}\ell_{1\text{split}}$
SSIM	0.712	0.741	0.752	<b>0.753</b>	0.739
NMSE	0.0925	<b>0.0646</b>	0.0655	0.0671	0.0719
PSNR	25.143	26.525	<b>26.535</b>	26.487	26.067

Further tuning to improve the ratio of high- to low- frequency led to better results for the higher cascade models. Notably, combining SSIM with the  $\perp\text{-}\ell_1$  loss was counter-productive for all metrics and suggests that further tuning of the weighting of each loss component is required.

**Introduction of the refinement module** Table 2 highlights the improvements in the quality of inference made by introducing the refinement module. Deploying the refinement as a separately trained post-processing module shows notable benefit improving more than adding an additional cascade to the plain CRNN. The end-to-end model results in further improvements upon separate training, of 4.4% in structural similarity and 3.9% in normalised mean square error relative to the plain CRNN, in spite of no longer being able to take advantage of the  $\perp$  loss.

**Table 2.** Performance comparisons of various model set-ups. Sequential (separate) and end-to-end training (combined) of the CRNN and refinement module are presented.

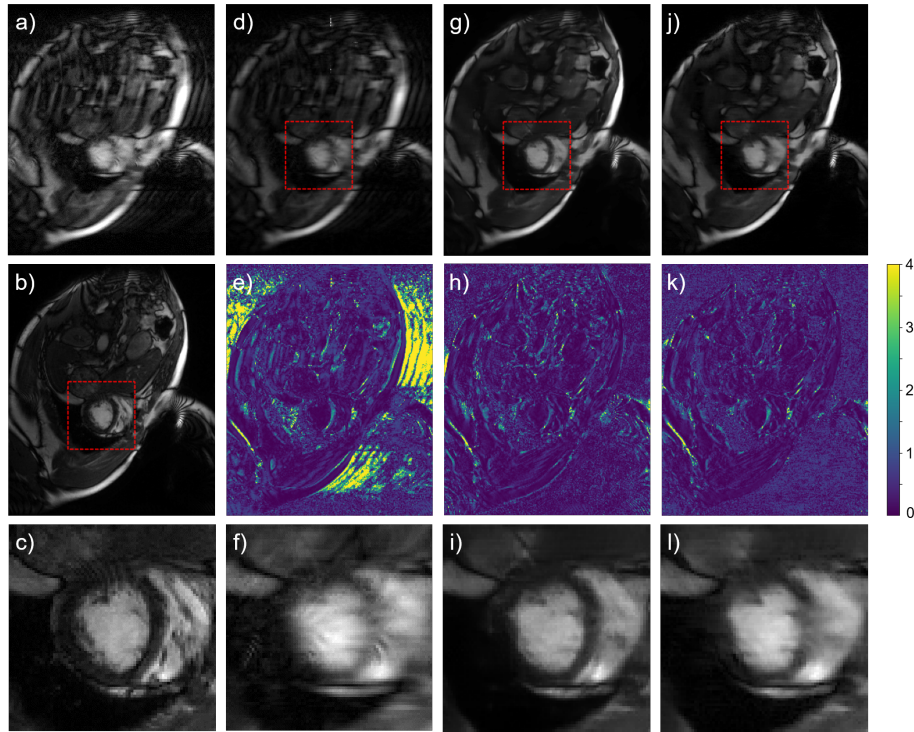
Cascades	6	6	6	7
Refinement module	None	Sequential	End-to-end	None
SSIM	0.768	0.792	<b>0.802</b>	0.765
NMSE	0.0516	0.0496	<b>0.0454</b>	0.0535
PSNR	27.354	27.597	<b>27.969</b>	27.351

Fig. 4 shows qualitatively the improvements made by the introduction of the refinement module at full scale. The error is reduced substantially and some finer details are resolved, though there is still scope for improvement at smaller scales. We generally see that the model is incapable of generating details that are completely lost in the undersampling process.

**Validation results** Our final tests prior to submission are presented in Table 3. Across all models, the short-axis reconstruction performs better quantitatively as there is more short-axis data available in training. For both views, the performance reduces with increased undersampling, as more detail is lost.

**Table 3.** Performance comparisons on evaluation of CMRxRecon cine cardiac MRI coil combined validation data for different acceleration rates (AR).

AR	Metric	Short-axis			Long-axis		
		U-Net Baseline	Plain CRNN	Proposed	U-Net Baseline	Plain CRNN	Proposed
4×	SSIM	0.641	0.824	<b>0.854</b>	0.573	0.757	<b>0.792</b>
	NMSE	0.180	0.0311	<b>0.0277</b>	0.188	0.0485	<b>0.0433</b>
	PSNR	23.084	29.842	<b>30.295</b>	22.040	26.958	<b>27.540</b>
8×	SSIM	0.637	0.796	<b>0.829</b>	0.574	0.723	<b>0.763</b>
	NMSE	0.201	0.0428	<b>0.0377</b>	0.191	0.0687	<b>0.0586</b>
	PSNR	22.603	28.364	<b>29.002</b>	22.234	25.588	<b>26.370</b>
10×	SSIM	0.641	0.788	<b>0.822</b>	0.588	0.717	<b>0.753</b>
	NMSE	0.210	0.0464	<b>0.0408</b>	0.198	0.0724	<b>0.064</b>
	PSNR	22.429	28.030	<b>28.644</b>	22.065	25.343	<b>25.965</b>

**Fig. 4.** Reconstruction (top) and associated error maps (middle) for the U-Net baseline and CRNN models. Finer details (bottom) are not resolved by the U-Net, which are partially captured by the plain CRNN model. The refinement module subsequently deblurs the image and provides better resolution at boundaries. (a)  $10 \times$  undersampled SAX input (b, c) fully sampled ground truth (d, e, f) U-Net (g, h, i) 6 cascades with combined refinement (j, k, l) 7 cascades, no refinement.



## 5 Discussion

**On model choice** Without extensive hyperparameterisation, the CRNN architecture demonstrated more promising performance than the CineNet, both qualitatively and quantitatively as shown in Fig. 2 and 3. In this network, the temporal average is subtracted from each slice and then the residuals are transformed into  $x-t$  and  $y-t$  axes being passed through a 2D U-Net structure. The 2D U-Net is computationally lightweight as the CineNet was originally designed for multi-coil radial acquisition processes, but even with an increased number of cascades fails to resolve finer details as clearly shown in Fig. 3. The recurrent connections exploit the temporal dependencies between slices more effectively than by attempting to capture these relationships through transformation and sparsification.

However, training of the plain CRNN has still resulted in lower perceptual quality than presented in the original implementation, which may have been further improved with more hyperparameterisation tuning. The large image size of up to  $246 \times 512$  proved challenging, and in our implementation limited us to 5 cascades of 48 channels for the original 24GB inference limitation for the CRNN network. In the original work [15], the model consisted of 10 cascades with 64 feature channels, which operated on smaller data with a smaller GPU.

**On the final model** The plain CRNN implementation substantially outperforms the baseline 2D U-Net, which has over double the number of trainable parameters, demonstrating the importance of exploiting temporal correlations. The reconstruction of the plain CRNN is a considerable improvement upon the  $10\times$  undersampled input, however smaller scale details that are resolved are blurry (e.g. Fig. 4l). The introduction of the low-cost refinement module led to better results with further denoising as presented in Fig. 4i. This shows promise for the implementation of lightweight single-image super-resolution models to assist in improving cardiac cine MRI reconstruction, either combined with the main reconstruction model or as a post-processing step. Relative to the ground truth, we still see that finer details are being missed that have been obscured due to the undersampling process. In [2], numerous MRI reconstruction experiments were conducted to test current deep learning reconstruction models. They proposed that “networks must be retrained on any subsampling patterns” for end to end CNN networks. In our approach, we trained all the acceleration rate together to get a stable but averaged reconstruction in a single model, subsequently resolving less finer details. The failure in generating details that have been lost, could be better tackled by a generative model [9]. As such, performance on patient volumes where more aliasing artefacts were present was poorer.

Therefore, to improve the proposed model, there is potential that training exclusively for each view can improve the final details. However, our model performs relatively well on the validation stage leaderboard for high acceleration factors, where finer details are more difficult to resolve, whilst we generally perform worse at lower acceleration factors. This suggests that whilst our model

failed to generate some finer details, other architectures also struggled once these details were lost or heavily obscured. There are numerous further modules that could have been implemented, had more time been available.

**On the loss function** We found that introduction of the  $\perp$  loss to the  $\ell_1$  loss improved both SSIM and PSNR, though at the expense of a slightly reduced  $\ell_1$  value itself. Treating the weightings of the loss functions as learnable parameters could lead to improved results in all metrics, as anticipated due to the results presented in [17]. Likewise, the introduction of the high-pass filter loss to focus the  $\ell_1$  loss on higher-frequency information increases the complexity of optimisation but was beneficial after the weightings were improved, though not presented quantitatively here.

## 6 Conclusions

In this challenge, we deployed a CRNN network combined with a refinement module to perform MRI reconstruction of cardiac cine data. We train the model for a range of acceleration factors and views using a high-pass filter to focus our loss on high-frequency details. From the quantitative analysis of the evaluation data and from direct viewing of the validation portion of the training data, the refinement module provides additional image quality with improvements of around 4% in all metrics relative to the plain CRNN implementation. As is typically found, some finer details at smaller scales remain unresolved that may be improved upon with further hyperparameter tuning and new modules. Nonetheless, the improvement upon the baseline is substantial and our model shows promise for improving cardiac MRI reconstruction.

## Acknowledgements

This work was supported in part by National Institutes of Health (NIH) grant 7R01HL148788-03. C. Jordan, Y. Du and Y. Xue thank additional financial support from the School of Engineering, the University of Edinburgh. S.A. Tsafaris also acknowledges the support of Canon Medical and the Royal Academy of Engineering and the Research Chairs and Senior Research Fellowships scheme (grant RCSR1819\8\25). The authors would like to thank Dr. Chen and K. Vilouras for inspirational discussions and assistance.

## References

1. Aggarwal, H.K., Mani, M.P., Jacob, M.: MoDL: Model-Based Deep Learning Architecture for Inverse Problems. *IEEE Transactions on Medical Imaging* **38**(2), 394–405 (2019). <https://doi.org/10.1109/TMI.2018.2865356>
2. Antun, V., Renna, F., Poon, C., Adcock, B., Hansen, A.C.: On instabilities of deep learning in image reconstruction and the potential costs of ai. *Proceedings of the National Academy of Sciences* **117**(48), 30088–30095 (2020)
3. Bilecen, B.B., Ayazoglu, M.: Bicubic++: Slim, Slimmer, Slimmest – Designing an Industry-Grade Super-Resolution Network (2023), <https://arxiv.org/abs/2305.02126>
4. CMRxRecon: Cardiac MRI Reconstruction Challenge. <https://cmrxrecon.github.io/>, Accessed: 26 Jul. 2023.
5. Hammernik, K., Klatzer, T., Kobler, E., Recht, M.P., Sodickson, D.K., Pock, T., Knoll, F.: Learning a variational network for reconstruction of accelerated MRI data. *Magnetic Resonance in Medicine* **79**(6), 3055–3071 (2018). <https://doi.org/10.1002/mrm.26977>
6. Han, X., Liu, Y., Lin, Y., Chen, K., Zhang, W., Liu, Q.: MDAMF: Reconstruction of Cardiac Cine MRI under Free-breathing using Motion-guided Deformable Alignment and Multi-resolution Fusion (2023), <https://arxiv.org/abs/2303.04968>
7. Han, Y., Yoo, J., Kim, H.H., Shin, H.J., Sung, K., Ye, J.C.: Deep learning with domain adaptation for accelerated projection-reconstruction MR. *Magnetic Resonance in Medicine* **80**(3), 1189–1205 (2018). <https://doi.org/10.1002/mrm.27106>
8. Hyun, C.M., Kim, H.P., Lee, S.M., Lee, S., Seo, J.K.: Deep learning for undersampled MRI reconstruction. *Physics in Medicine & Biology* **63**(13), 135007 (jun 2018). <https://doi.org/10.1088/1361-6560/aac71a>, <https://dx.doi.org/10.1088/1361-6560/aac71a>
9. Jalal, A., Arvinte, M., Daras, G., Price, E., Dimakis, A.G., Tamir, J.: Robust compressed sensing mri with deep generative priors. *Advances in Neural Information Processing Systems* **34**, 14938–14954 (2021)
10. Kofler, A., Haltmeier, M., Schaeffter, T., Kolbitsch, C.: An end-to-end-trainable iterative network architecture for accelerated radial multi-coil 2D cine MR image reconstruction. *Medical Physics* **48**(5), 2412–2425 (2021). <https://doi.org/10.1002/mp.14809>
11. Küstner, T., Fuin, N., Hammernik, K., Bustin, A., Qi, H., Hajhosseiny, R., Masci, P.G., Neji, R., Rueckert, D., Botnar, R.M., Prieto, C.: CINENet: deep learning-based 3D cardiac CINE MRI reconstruction with multi-coil complex-valued 4D spatio-temporal convolutions. *Scientific Reports* **10**(13710) (2020). <https://doi.org/10.1038/s41598-020-70551-8>
12. Lee, D., Yoo, J., Tak, S., Ye, J.C.: Deep Residual Learning for Accelerated MRI Using Magnitude and Phase Networks. *IEEE Transactions on Biomedical Engineering* **65**(9), 1985–1995 (2018). <https://doi.org/10.1109/TBME.2018.2821699>
13. Lyu, J., Li, G., Wang, C., Qin, C., Wang, S., Dou, Q., Qin, J.: Region-focused multi-view transformer-based generative adversarial network for cardiac cine MRI reconstruction. *Medical Image Analysis* **85**, 102760 (2023). <https://doi.org/10.1016/j.media.2023.102760>
14. Patel, D., Sastry, P.S.: Memorization in Deep Neural Networks: Does the Loss Function Matter? In: *Advances in Knowledge Discovery and Data Mining*, pp. 131–142. Springer International Publishing, Cham (2021). [https://doi.org/10.1007/978-3-030-75765-6\\_11](https://doi.org/10.1007/978-3-030-75765-6_11)

15. Qin, C., Schlemper, J., Caballero, J., Price, A.N., Hajnal, J.V., Rueckert, D.: Convolutional Recurrent Neural Networks for Dynamic MR Image Reconstruction. *IEEE Transactions on Medical Imaging* **38**(1), 280–290 (2019). <https://doi.org/10.1109/TMI.2018.2863670>
16. Schlemper, J., Caballero, J., Hajnal, J.V., Price, A.N., Rueckert, D.: A Deep Cascade of Convolutional Neural Networks for Dynamic MR Image Reconstruction. *IEEE Transactions on Medical Imaging* **37**(2), 491–503 (2018). <https://doi.org/10.1109/TMI.2017.2760978>
17. Terpstra, M.L., Maspero, M., Sbrizzi, A., van den Berg, C.A.:  $\perp$ -loss: A symmetric loss function for magnetic resonance imaging reconstruction and image registration with deep learning. *Medical Image Analysis* **80**, 102509 (2022). <https://doi.org/10.1016/j.media.2022.102509>
18. Tezcan, K.C., Baumgartner, C.F., Luechinger, R., Pruessmann, K.P., Konukoglu, E.: MR Image Reconstruction Using Deep Density Priors. *IEEE Transactions on Medical Imaging* **38**(7), 1633–1642 (2019). <https://doi.org/10.1109/TMI.2018.2887072>
19. Vornehm, M., Wetzl, J., Giese, D., Ahmad, R., Knoll, F.: Spatiotemporal variational neural network for reconstruction of highly accelerated cardiac cine MRI. *European Heart Journal - Cardiovascular Imaging* **23**(Supplement 2), 34–35 (08 2022). <https://doi.org/10.1093/ehjci/jeac141.018>
20. Wang, C., Li, Y., Lv, J., Jin, J., Hu, X., Kuang, X., Chen, W., Wang, H.: Recommendation for Cardiac Magnetic Resonance Imaging-Based Phenotypic Study: Imaging Part. *Phenomics* **1**, 151–170 (2021). <https://doi.org/10.1007/s43657-021-00018-x>
21. 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 (2016). <https://doi.org/10.1109/ISBI.2016.7493320>
22. 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 (2018). <https://doi.org/10.1109/TMI.2017.2785879>
23. Yang, J., Küstner, T., Hu, P., Liò, P., Qi, H.: End-to-End Deep Learning of Non-rigid Groupwise Registration and Reconstruction of Dynamic MRI. *Frontiers in Cardiovascular Medicine* **9** (2022). <https://doi.org/10.3389/fcvm.2022.880186>
24. Zhang, T., Pauly, J.M., Vasanawala, S.S., Lustig, M.: Coil compression for accelerated imaging with Cartesian sampling. *Magnetic Resonance in Medicine* **69**(2), 571–582 (2013). <https://doi.org/https://doi.org/10.1002/mrm.24267>
25. Zhang, Y., Hu, Y.: Dynamic Cardiac MRI Reconstruction Using Combined Tensor Nuclear Norm and Casorati Matrix Nuclear Norm Regularizations. In: 2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI). pp. 1–4 (2022). <https://doi.org/10.1109/ISBI52829.2022.9761409>