

---

# Self-supervised Low-rank plus Sparse Network for Radial MRI Reconstruction

---

Andrei Mancu<sup>1</sup>, Wenqi Huang<sup>2</sup>, Gastao Lima da Cruz<sup>3</sup>, Daniel Rueckert<sup>2,4</sup>, and Kerstin Hammernik<sup>1</sup>

<sup>1</sup>School of Computation, Information and Technology, Technical University of Munich

<sup>2</sup>Klinikum Rechts der Isar, Technical University of Munich

<sup>3</sup>Department of Radiology, University of Michigan

<sup>4</sup>Department of Computing, Imperial College London

## Abstract

In this work, we introduce a physics-guided self-supervised learning approach to reconstruct dynamic magnetic resonance (MR) images from sparsely sampled radial cardiac data. The architecture incorporates a variable splitting scheme via a quadratic penalty approach consisting of iterative data consistency and denoiser steps. To accommodate cardiac motion, the denoiser implements a learnable low-rank and sparse component instead of a conventional convolutional neural network. We compare the proposed model to iterative regularized MRI reconstruction techniques and to other deep neural network approaches adapted to radial data, both in supervised and self-supervised tasks. Our proposed method surpasses the performance of other techniques for a single heartbeat and four heartbeat MR image reconstruction. Furthermore, our approach outperforms other deep neural network reconstruction approaches in both supervision and self-supervision tasks.

## 1 Introduction

Magnetic Resonance Imaging (MRI) depicts anatomical structures and functional processes within the human body in high resolution, which makes MRI a leading diagnostic tool in medical imaging. To reduce scan time, fewer measurements in signal space (k-space) are acquired (1), often called undersampling of the acquisition process. However, undersampling the k-space leads to an ill-posed image reconstruction problem, resulting in aliasing artifacts and noise enhancement (2; 3). Parallel Imaging (4; 5) and Compressed Sensing (6) mitigate these issues, however, they face challenges at higher undersampling rates.

Recently, physics-guided Deep Neural Networks (DNNs) (7; 8; 9) have been introduced to withstand these artifacts. However, these approaches typically require lots of training data to operate effectively. Recent advancements in self-supervised learning (ssl) have allowed training only on undersampled datasets (10). One example was proposed by Yaman et al. (11), allowing one to reconstruct MR images from a single subject using a zero-shot self-supervised learning approach.

However, two key concerns persist in current MRI self-supervision approaches. Firstly, Yaman et al.’s network faces challenges in accurately capturing moving organs due to its inability to leverage the intrinsic properties of dynamic MR images, such as a strong low-rank prior (12; 13; 14). Huang et al. (15) demonstrated that by extracting the low-rank component and learning the sparse component of an image, undersampling artifacts can be substantially reduced in dynamic MRI reconstruction. Secondly, while Cartesian acquisitions in Yaman et al.’s work are susceptible to motion artifacts, radially sampled MRI datasets (16; 17) are more robust (18). As multiple radial spokes pass the low-frequency k-space, motion-corrupted spokes are compensated through averaging (18). Additionally,

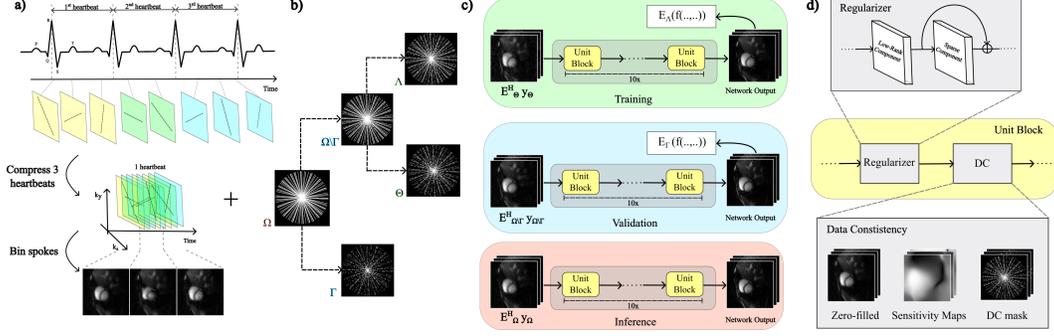


Figure 1: a) Data Processing. Sparse measurements from three heartbeats are compressed into one heartbeat and binned into temporal windows to create three cardiac phases. b) Radial Mask Splitting for Self-Supervision. The initial mask  $\Omega$  is split into a validation mask  $\Gamma$  and a remainder mask  $\Omega \setminus \Gamma$ . The remainder mask is randomly split into a training mask  $\Theta$  and validation mask  $\Lambda$  at each training step. c) Model training, validation and inference. Training input is created by masking the sample’s k-space with mask  $\Theta$ . Validation input is masked with  $\Omega \setminus \Gamma$  and inference input with  $\Omega$ . d) Components of an Unrolled Network Unit Block. The regularizer includes low-rank and sparse components, while data consistency uses a classic conjugate gradient approach.

golden-angle radial sampling provides non-overlapping k-space coverage over time, allowing for flexible retrospective data binning (18).

In our work, we introduce a self-supervised DNN for reconstructing cardiac MR images from radially sampled data. Our model adopts a physics-guided deep learning reconstruction approach, enhancing its robustness to artifacts and noise (7; 19). The model employs iterative algorithm unrolling, alternating between data consistency (DC) and regularization steps (20; 21). In place of a conventional denoiser, our approach utilizes learnable low-rank and sparse components as regularizers, effectively capturing organ motion.

Our key contributions include: 1) Pioneering the use of radially acquired MRI measurements for self-supervision, allowing models to learn from limited radial data, undersampled scans, and exploiting advantages of radial data acquisition. 2) Enhancing the architecture to accommodate cardiac motion by incorporating a low-rank and sparse denoiser, substantially improving MRI reconstruction for moving organs.

## 2 Theoretical Background and Methods

### 2.1 Data Processing and Undersampled MRI

Cardiac MRI measurements are acquired simultaneously with the patient’s electrocardiogram (ECG) signal, which can cover multiple cardiac cycles. To address the sparsity of radial measurements in the spatial-frequency dimensions, measurements are aligned into a single cardiac cycle and then grouped into temporal windows, each corresponding to a specific cardiac phase in the heartbeat. Figure 1a) illustrates the data binning for three heartbeats.

Current MRI systems simultaneously acquire frequency signals from multiple receiver coils, each sensitive to different regions of the scanned subject. These signals are then combined to create a unified image space. A typical representation of a multi-coil MRI forward model is given by:

$$y_{\Omega} = E_{\Omega}x + n, \quad E_{\Omega} := P_{\Omega}FC \quad (1)$$

where  $x$  is the target image,  $y_{\Omega}$  is the k-space measurements on an undersampling pattern  $\Omega$ , and  $n$  is the noise (4).  $E_{\Omega}$  models the forward operator comprised of the undersampling operator  $P_{\Omega}$ , the Fourier transform  $F$ , and the coil sensitivity maps  $C$ .

To reconstruct the image  $x$  from the frequency measurements  $y_{\Omega}$ , one must solve the regularized inverse problem given by

$$\underset{x}{\operatorname{argmin}} \|y_{\Omega} - E_{\Omega}x\|_2^2 + R(x). \quad (2)$$

Here, the consistency (DC) term  $\|y_\Omega - E_\Omega x\|_2^2$  aims to minimize the discrepancy between the reconstructed image and the measured k-space data, while the regularization term  $R(x)$  is employed to control noise and maintain image quality.

## 2.2 Learned physics-guided L+S reconstruction

Physics-guided deep learning reconstruction is one category of deep learning methods designed to solve the inverse problem in Eq. 2 by two sub-problems via variable-splitting Eq. 2:

$$z^{(i)} = f(y_\Omega^n, E_\Omega^n; \phi), \quad (3a)$$

$$x^{(i)} = x^{(i-1)} - \mu E_\Omega^H (E_\Omega z^{(i)} - y_\Omega), \quad (3b)$$

where  $z^{(i)}$  is an intermediate variable,  $x^{(i)}$  the reconstructed image at iteration ( $i$ ) within a fixed number of iterations in the unrolled algorithm (21),  $\mu$  is a regularization parameter, and  $f$  is the denoiser function parameterized by  $\phi$ . Usually, Eq. 3b, representing the DC sub-problem, is solved with a linear optimizer based on either gradient descent methods (7) or conjugate gradient methods (22).

As for Eq. 3a, where conventionally a DNN is responsible for the regularization, our model integrates L+S (15), which incorporates a low-rank component  $l$  for adaptively thresholding singular values (23) of the input image  $x$  and a sparse component  $s$  to capture motion through a few convolutional layers. This transforms Eq. 3 into:

$$l^{(i)}, s^{(i)} = f(y_\Omega^n, E_\Omega^n; \phi), \quad (4a)$$

$$x^{(i)} = x^{(i-1)} - \mu E_\Omega^H (E_\Omega (l^{(i)} + s^{(i)}) - y_\Omega). \quad (4b)$$

**Supervised and Self-Supervised Learning** We adopt an extended *self-supervised* approach to (11) to radially sampled dynamic MRI data in our work.  $\Omega$  represents the whole binned sample of spokes. We subsample  $\Omega$  to generate the validation mask  $\Gamma$  and a remainder mask  $\Omega \setminus \Gamma$ . Thus, the model parameters can be optimized by minimizing the loss in k-space for the model parameters  $\phi$ :

$$\min_{\phi} \frac{1}{N} \sum_{n=1}^N L(y_\Lambda^n, E_\Lambda^n (f(y_\Theta^n, E_\Theta^n; \phi))), \quad (5)$$

where  $y_\Lambda^n$  and  $y_\Theta^n$  are the  $n^{th}$  sample's k-space masked with  $\Lambda$  and  $\Theta$ , respectively.  $E_\Lambda^n$  and  $E_\Theta^n$  are the encoding operators for the  $n^{th}$  sample, utilizing  $\Lambda$  and  $\Theta$  masks. The model also supports *supervised* learning 5.3, by replacing every mask with the initial undersampled k-space mask  $\Omega$  and calculating the loss in the image domain. Figure 1b and Figure 1c depict a complete overview of the self-supervised masking.

## 2.3 Experimental Setting

We compare the proposed model to state-of-the-art linear reconstruction (CGSense (4)), compressed sensing reconstruction (L+S (15)) and deep learning reconstruction model proposed by Yaman et al. (11), where we consider both supervised and self-supervised training. The CGSense algorithm (4) was applied with a maximum iteration number of 6. The L+S algorithm (14) was used with 50 maximum iterations, low-rank regularization  $\lambda_L = 0.0037$ , sparse regularization  $\lambda_S = 0.1$  and step size gamma  $\gamma = 1$ . Our proposed model in Eq. 4 as well as the model by Yaman et al. in Eq. 3, use 10 iterations of their unrolled blocks (Figure 1d). Yaman et al. (11) uses a 5-block ResNet as the regularizer with 64 features maps in each block. Our network has 3 convolutional layers of 32 filters and leakyReLU activations (24) for the sparse component and a learnable low-rank regularizing parameter  $\lambda_L$  for each unroll block. Both architectures employ conjugate gradient for data consistency. For both self-supervised models, the validation mask  $\Gamma$  is subsampled 20% from the initial  $\Omega$  and a varying 60% – 40%  $\Theta$  and  $\Lambda$  masks from the remaining, fixed  $\Omega \setminus \Gamma$  mask. The loss of the models in the supervised case is the MSE loss (25), while the self-supervised loss is a combination of normalized mean squared error and absolute error as used in (11).

The training and evaluation is conducted on 128 slices from 16 subjects, for a single cardiac cycle and for four cardiac cycles. Additional information on the data can be found in the Supplementary Material 5. The evaluation metrics are normalized mean squared error (NMSE) (25), structural similarity (SSIM) (26) and peak signal to noise ratio (PSNR) (27).

Table 1: Inference results of all experiments. The mean and the standard deviations of NMSE, SSIM and PSNR are displayed for each model prediction on 1 or 4 cardiac cycle dataset.

Cardiac Cycles	Metric	CGSense	L+S	Yaman et al. (s)	Proposed (s)	Yaman et al. (SSL)	Proposed (SSL)
1	NMSE	0.18±0.05	0.10±0.04	0.12±0.04	0.12±0.05	0.12±0.05	<b>0.09±0.04</b>
	SSIM	0.65±0.04	0.78±0.06	0.78±0.05	0.74±0.05	0.77±0.05	<b>0.83±0.05</b>
	PSNR	26.91±1.84	29.83±2.47	28.95±2.32	29.41±2.20	28.74±2.31	<b>31.55±2.71</b>
4	NMSE	0.08±0.05	0.04±0.04	0.05±0.03	0.04±0.03	0.06±0.05	<b>0.03±0.02</b>
	SSIM	0.76±0.048	0.90±0.05	0.87±0.04	0.91±0.04	0.86±0.05	<b>0.93±0.03</b>
	PSNR	31.45±2.34	35.50±3.37	33.45±2.53	35.22±2.92	33.00±2.82	<b>36.12±2.93</b>

### 3 Results

Table 1 presents metrics for both test sets, providing mean and standard deviation values. Our proposed method outperforms both iterative reconstruction methods (CGSense and L+S), and Yaman et al.’s approach in all metrics. Notably, L+S algorithm exhibits superior PSNR than techniques without the low-rank prior component of cardiac images, while performing similarly in NMSE and SSIM. Furthermore, self-supervision variants of both Yaman et al. and our method perform on par with or even outperform the supervised variants.

Compared to other methods, our proposed method effectively reduced streaking artifacts and enhanced contrast, even at high undersampling rates, as shown in Figure 2. More detailed error maps comparisons can be found in Supplementary Materials 5.

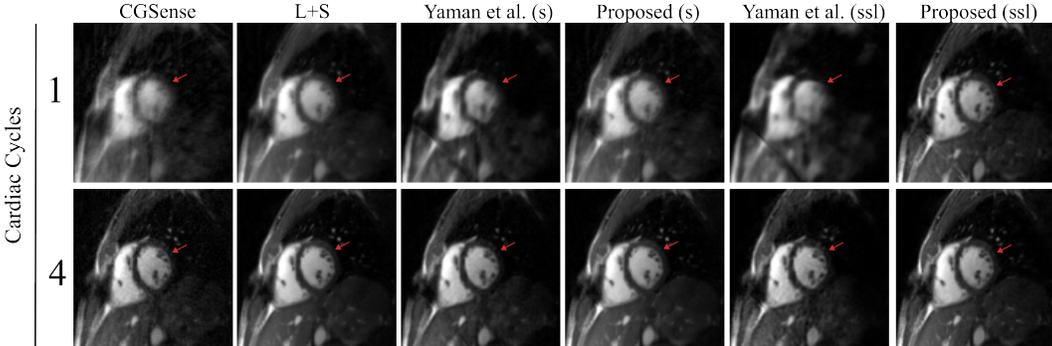


Figure 2: Prediction images on a sample from the test set. First row presents the 1 cardiac cycle dataset, while the second row presents the 4 cycle dataset. "s" stands for *supervised*. The arrows are pointing at the contrast of the myocardium of the heart.

### 4 Discussion and Conclusion

Incorporating the low-rank and sparse properties into our model has proven highly advantageous for reconstruction, particularly in dynamic data scenarios, where Yaman et al.’s approach (11) faces challenges in both supervised and self-supervised variants. Furthermore, adapting L+S into our approach yields additional benefits, as the model has approximately three times less trainable parameters than (11), resulting in reduced mathematical complexity, faster training, and inference.

The self-supervised version of our approach outperforms the supervised version in all metrics. This phenomenon may arise from the repeated application of the non-linear Fast Fourier Transform, causing fluctuations across iterations in the unrolled network. When combined with the MSE loss, which is sensitive to image intensity outliers, it may hinder the model’s training.

In conclusion, we developed a self-supervised physics-guided model tailored for radially sampled cardiac MR data, building on (11). Our model leverages the low-rank and sparse properties of moving organs to handle dynamic data more effectively, especially in a limited data scenario.

## References

- [1] J. Tsao and S. Kozerke, "MRI temporal acceleration techniques," *Journal of Magnetic Resonance Imaging*, vol. 36, no. 3, pp. 543–560, 2012.
- [2] G. Cruz, B. Yaman, M. Akçakaya, M. Doneva, and C. Prieto, "Chapter 2 - MRI Reconstruction as an Inverse Problem," in *Magnetic Resonance Image Reconstruction*, vol. 7 of *Advances in Magnetic Resonance Technology and Applications*, pp. 37–57, Academic Press, 2022.
- [3] C. Sandino, J. Cheng, F. Chen, M. Mardani, J. Pauly, and S. Vasanawala, "Compressed Sensing: From Research to Clinical Practice With Deep Neural Networks: Shortening Scan Times for Magnetic Resonance Imaging," *IEEE Signal Processing Magazine*, vol. 37, pp. 117–127, 01 2020.
- [4] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, "Sense: Sensitivity encoding for fast MRI," *Magnetic Resonance in Medicine*, vol. 42, no. 5, pp. 952–962, 1999.
- [5] M. Griswold, P. Jakob, R. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase, "Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA)," *Magnetic Resonance in Medicine*, vol. 47, pp. 1202–10, 06 2002.
- [6] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse mri: The application of compressed sensing for rapid mr imaging," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.
- [7] K. Hammernik, T. Klatzer, E. Kobler, M. Recht, D. Sodickson, T. Pock, and F. Knoll, "Learning a Variational Network for Reconstruction of Accelerated MRI Data," *Magnetic Resonance in Medicine*, pp. 3055–3071, 2017.
- [8] F. Knoll, K. Hammernik, C. Zhang, S. Moeller, T. Pock, D. Sodickson, and M. Akçakaya, "Deep-Learning Methods for Parallel Magnetic Resonance Imaging Reconstruction: A survey of the current approaches, trends, and issues," *IEEE Signal Processing Magazine*, vol. 37, pp. 128–140, 01 2020.
- [9] M. Akçakaya, B. Yaman, H. Chung, and J. C. Ye, "Unsupervised Deep Learning Methods for Biological Image Reconstruction and Enhancement: An overview from a signal processing perspective," *IEEE Signal Processing Magazine*, vol. 39, no. 2, pp. 28–44, 2022.
- [10] B. Yaman, S. A. H. Hosseini, S. Moeller, J. Ellermann, K. Uğurbil, and M. Akçakaya, "Self-supervised learning of physics-guided reconstruction neural networks without fully sampled reference data," *Magnetic Resonance in Medicine*, vol. 84, no. 6, pp. 3172–3191, 2020.
- [11] B. Yaman, S. A. H. Hosseini, and M. Akçakaya, "Zero-Shot Self-Supervised Learning for MRI Reconstruction," in *International Conference on Learning Representations*, 2022.
- [12] P. J. Shin, P. E. Z. Larson, M. A. Ohliger, M. Elad, J. M. Pauly, D. B. Vigneron, and M. Lustig, "Calibrationless parallel imaging reconstruction based on structured low-rank matrix completion," *Magnetic Resonance in Medicine*, vol. 72, no. 4, pp. 959–970, 2014.
- [13] H. Gao, J.-F. Cai, Z. Shen, and H. Zhao, "Robust principal component analysis-based four-dimensional computed tomography," *Physics in Medicine & Biology*, vol. 56, no. 11, p. 3181, 2011.
- [14] R. Otazo, E. Candès, and D. K. Sodickson, "Low-rank plus sparse matrix decomposition for accelerated dynamic MRI with separation of background and dynamic components," *Magnetic Resonance in Medicine*, vol. 73, no. 3, pp. 1125–1136, 2015.
- [15] W. Huang, Z. Ke, Z.-X. Cui, J. Cheng, Z. Qiu, S. Jia, L. Ying, Y. Zhu, and D. Liang, "Deep low-rank plus sparse network for dynamic mr imaging," *Medical Image Analysis*, vol. 73, p. 102190, 2021.
- [16] K. T. Block, H. Chandarana, S. Milla, M. Bruno, T. Mulholland, G. Fatterpekar, M. Hagiwara, R. Grimm, C. Geppert, B. Kiefer, *et al.*, "Towards routine clinical use of radial stack-of-stars 3D gradient-echo sequences for reducing motion sensitivity," *Journal of the Korean Society of Magnetic Resonance in Medicine*, vol. 18, no. 2, pp. 87–106, 2014.

- [17] R. M. Azevedo, R. O. P. de Campos, M. Ramalho, V. Herédia, B. M. Dale, and R. C. Semelka, "Free-Breathing 3D T1-Weighted Gradient-Echo Sequence With Radial Data sampling in abdominal MRI: Preliminary Observations," *American Journal of Roentgenology*, vol. 197, no. 3, pp. 650–657, 2011. PMID: 21862807.
- [18] L. Feng, "Golden-Angle Radial MRI: Basics, Advances, and Applications," *Journal of Magnetic Resonance Imaging*, vol. 56, 04 2022.
- [19] S. A. H. Hosseini, B. Yaman, S. Moeller, M. Hong, and M. Akçakaya, "Dense Recurrent Neural Networks for Accelerated MRI: History-Cognizant Unrolling of Optimization Algorithms," *IEEE Journal of Selected Topics in Signal Processing*, vol. 14, no. 6, pp. 1280–1291, 2020.
- [20] J. Schlemper, J. Caballero, J. Hajnal, A. Price, and D. Rueckert, "A deep cascade of convolutional neural networks for dynamic mr image reconstruction," *IEEE Transactions on Medical Imaging*, vol. 37, pp. 491–503, Feb. 2018.
- [21] V. Monga, Y. Li, and Y. C. Eldar, "Algorithm Unrolling: Interpretable, Efficient Deep Learning for Signal and Image Processing," *IEEE Signal Processing Magazine*, vol. 38, no. 2, pp. 18–44, 2021.
- [22] H. K. Aggarwal, M. P. Mani, and M. Jacob, "MoDL: Model-Based Deep Learning Architecture for Inverse Problems," *IEEE Transactions on Medical Imaging*, vol. 38, no. 2, pp. 394–405, 2019.
- [23] V. Klema and A. Laub, "The singular value decomposition: Its computation and some applications," *IEEE Transactions on Automatic Control*, vol. 25, no. 2, pp. 164–176, 1980.
- [24] M. Khalid, J. Baber, M. K. Kasi, M. Bakhtyar, V. Devi, and N. Sheikh, "Empirical Evaluation of Activation Functions in Deep Convolution Neural Network for Facial Expression Recognition," in *2020 43rd International Conference on Telecommunications and Signal Processing (TSP)*, pp. 204–207, 2020.
- [25] P. Bickel and K. Doksum, *Mathematical Statistics: Basic Ideas and Selected Topics*. No. 1 in *Mathematical Statistics: Basic Ideas and Selected Topics*, Prentice Hall, 2001.
- [26] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [27] O. S. Faragallah, H. El-Hoseny, W. El-Shafai, W. A. El-Rahman, H. S. El-Sayed, E.-S. M. El-Rabaie, F. E. A. El-Samie, and G. G. N. Geweid, "A Comprehensive Survey Analysis for Present Solutions of Medical Image Fusion and Future Directions," *IEEE Access*, vol. 9, pp. 11358–11371, 2021.

## 5 Supplementary Material

In this section we provide additional details about the setup of data, the model and the experiments that we have conducted. Then, the model predictions are analyzed more in-depth by providing error maps time profiles of the predictions.

### 5.1 Data

Data was acquired on a 1.5T Philips (Ingenia, Best, Netherlands) machine using a 28-channel cardiac coil. The privately acquired data consists of 8960 radial spokes and ECG signals of 128 samples from 16 anonymised volunteer subjects. Slices were acquired on the short axis, with a field of view of  $256 \times 256 \text{ mm}^2$ , 8 mm slice thickness,  $2 \times 2 \text{ mm}^2$  resolution, TE/TR = 1.16/2.3 ms, b-SSFP readout, radial tiny golden angle of  $23.6^\circ$ , flip angle  $60^\circ$ , 20s nominal scan time and breath-hold acquisition.

After binning, each sample in the dataset consists of 30 cardiac phases (30 images) corresponding to the feature dimension of the model input. Depending on the heart rate, the number of acquired spokes per cardiac phase ranges from 11 to 19 spokes for the one heartbeat dataset and 45 to 74 spokes for the four heartbeat dataset. We normalize the data such that the operator norm of the non-linear Fast Fourier Transform becomes 1.

### 5.2 Model Regularizer

Subsection 2.2 presented the regularizer and data consistency steps within a unit block of the model architecture. Figure 3 provides an in-depth representation of the regularizer, which contains the low-rank component with a learnable parameter that thresholds the singular values of the dynamic input image. Consequently, only the non-moving parts of the image over time are retained, such as the liver and ribs. Subsequently, the sparse component eliminates streaking artifacts and improves contrast through three 3D convolutions. The sparse component is responsible for the moving parts of the image, such as the heart's beating action.

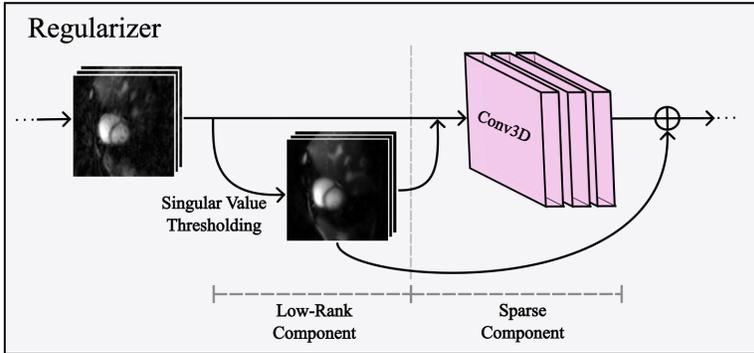


Figure 3: An in-depth representation of the regularizer workflow on a sample from the dataset. The low-rank component reconstructs the input images with thresholded singular values. The low-rank images are then concatenated with the input images and run through 3 convolutional layers with leakyReLU activations. The output of the sparse component is added to the low-rank images to generate the output of the regularizer.

### 5.3 Supervised Learning

Subsection 2.2 presented the self-supervised loss for network training. Here, we provide the supervised variant that we used in the experiments. This variant only uses the fully undersampled mask  $\Omega$  and calculates the loss in the image space with regard to the reference 20 heartbeat L+S reconstruction. This is motivated by the fact that, after spoke binning, we lose the target k-space in its original form, to which we could compare.

$$\min_{\phi} \frac{1}{N} \sum_{n=1}^N L(x_{ref}^n, f(y_{\Omega}^n, E_{\Omega}^n; \phi)), \quad (6)$$

where  $x_{ref}$  is the target image and  $E_{\Omega}$  the encoding operator using the fully undersampled k-space mask  $\Omega$ .

#### 5.4 Output Analysis

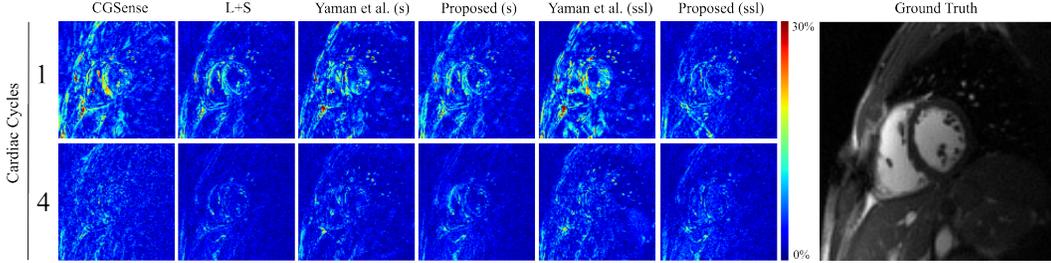


Figure 4: Error maps of the prediction images on a sample from the test set. First row presents the 1 cardiac cycle dataset, while the second row presents the 4 cycle dataset. The value scale represents the percentage of intensity error with regard to the ground truth maximum pixel intensity.

Figure 4 shows the error maps of the image predictions presented in Figure 2. The error maps show that our approach reduces the amount of streaking artifacts and noise compared to other methods, while still preserving structural integrity, even for high undersampling rates. Indeed, it can be seen that the methods implementing the low-rank prior present less error than the techniques that do not.

Specifically, the one heartbeat reconstructions for L+S and our proposed model contain few intensity errors inside and outside the heart. Our proposed method achieves the least error on the heart’s outer walls. Conversely, CGSense and Yaman et al.’s approach not only contain more noise outside the heart, but present high intensity errors inside of the heart’s region.

In the case of the four heartbeat reconstructions, all methods achieve low error. However, the advantages of the low-rank and sparse components can also be seen here. The methods not implementing the two components show higher noise than the ones that do so. Thus, we confirm that the low-rank prior approach delivers improved results on moving MRI scans.

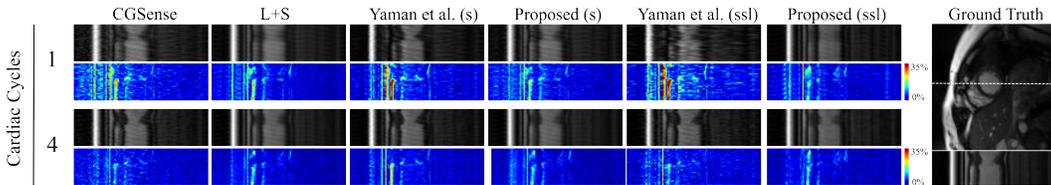


Figure 5: Time profile (yt plane) of the predictions on a sample from the two datasets (one and four heartbeat reconstructions). The value scale of the error maps represents the percentage of intensity error with regard to the ground truth maximum pixel intensity.

Figure 5 shows the time profiles in the yt plane of the predictions from a sample of the dataset. Compared to all other reconstructions, the proposed model manages to keep the error low, while still preserving good structural integrity of the moving heart, even for a single cardiac cycle reconstruction. Furthermore, in the case of the single cardiac cycle reconstruction, the L+S algorithm and our proposed method present the least amount of total intensity error, especially around the heart’s region. Yaman et al.’s approach struggles to reconstruct the tissue outside the heart, presenting a high error rate at the chest walls of the patient. In the four cardiac cycle case, our technique and L+S achieve the lowest error, while Yaman et. al’s models present the highest amount of noise. Moreover, Yaman et al.’s models and CGSense struggle to find the right intensities for the regions outside the heart, a problem that is solved in the other models by the low-rank component. This enforces again the importance of the low-rank prior and the learnable sparse component.