

PhaseGen: A Diffusion-Based Approach for Complex-Valued MRI Data Generation

Moritz Rempe^{1,2}[0000-0002-7137-2624], Fabian Hörst¹[0000-0002-8906-7644],
Helmut Becker¹, Marco Schlimbach⁴[0009-0003-6664-5848], Lukas T.
Rotkopf^{1,7}, Kevin Kröninger⁴[0000-0001-9873-0228], and Jens Kleesiek^{1,2,3,4,5,6}

¹ Institute for AI in Medicine (IKIM), University Hospital Essen, 45147 Essen, Germany

² Cancer Research Center Cologne Essen (CCCE), University Medicine Essen, 45147 Essen, Germany

³ RACOON Study Group, Site Essen, Essen Germany

⁴ Department of Physics, Technical University Dortmund, 44227 Dortmund, Germany

⁵ German Cancer Consortium (DKTK), Partner Site Essen, 45147 Essen, Germany

⁶ Medical Faculty and Faculty of Computer Science, University of Duisburg-Essen, 45141 Essen, Germany

⁷ German Cancer Research Center (DKFZ), 69120 Heidelberg, Germany
moritz.rempe@uk-essen.de

Abstract. Magnetic resonance imaging (MRI) raw data (k-Space) contains valuable magnitude and phase information, yet clinical AI methods typically discard phase data. To address this, we introduce a magnitude-conditioned generative phase modeling framework that enables synthesis of complex-valued k-Space data from widely available magnitude-only MRI datasets. Pretraining with PhaseGen significantly improves k-Space skullstripping accuracy (from 40.1% to 80.1%) and enhances MRI reconstruction. This approach successfully bridges the gap between abundant magnitude-only datasets and information-rich complex-valued MRI data, enabling more accurate diagnostic tasks.

Keywords: MRI Raw Data · GenAI · Complex-valued neural networks.

1 Introduction

Magnetic resonance imaging (MRI) is one of the most common clinical imaging procedures. Due to its high resolution and ability to visualize soft tissue, MRI is used in a wide range of medical applications, including the diagnosis of cancer, neurological disorders, and musculoskeletal diseases. With the advent of artificial intelligence (AI), new methods for improving or automating these medical applications are being developed [1–3]. While most of these methods are designed for the usage of magnitude image domain data [4], the initial MRI raw data is acquired in the so-called "k-Space", the frequency domain. This




Fig. 1. Overview of the proposed method. Publicly available magnitude image domain data is used to generate synthetic complex-valued k-Space data. This synthetic data facilitates pretraining of models for clinical downstream tasks, which can later be fine-tuned using real-world data.

raw data is complex-valued, including magnitude and phase data. While this frequency domain data is not interpretable by humans, this additional data - when compared to the commonly used magnitude data in the image domain - provides neural networks with additional information. Until now, this raw data has primarily been used for reconstruction and transformation into the image domain and discarded afterwards. While recent work has demonstrated the benefits of using these complex-valued raw data in diagnostic tasks [5, 6], research and publicly available datasets are scarce. Dishner et al. state a total of 110 MRI datasets containing 1.68 million individual subjects, with potential use for AI [7]. This vast amount of data only includes the magnitude image domain data, with the additional phase information of complex-valued raw data being discarded. When looking at the available MRI k-Space datasets, there is only a handful of contributions, such as the popular "FastMRI" datasets [8–10], with roughly 9300 subjects in total. Moreover, it should be noted that the publicly available raw data is exclusively intended for the purpose of reconstructing undersampled MRI data and so far not for diagnostic purposes. In contrast to image domain datasets, it is not intended for tasks such as classification or segmentation [11]. Due to the lack of available data, or necessary labels for classification or segmentation tasks, there is a lot of untapped potential of MRI research. In the past few years, especially with the rise of diffusion models, the field of generative AI experienced a lot of attention. With the help of generative models, synthetic datasets can be created which improve model performance, by artificially expanding the available training data [12]. In this work, we propose a magnitude-conditioned generative phase modeling framework that enables the synthesis of complex-valued MRI raw data from magnitude-only datasets. We demonstrate that such synthetic phase modeling improves downstream task performance and significantly reduces the amount of real phase data required. This approach comes with the benefit of reducing the need for real-world data, which is often scarce and hard to obtain. Focus can then be on less, but high quality data, which can be used for fine-tuning. As an example, Fig. 1 shows the workflow for a possi-

ble application of our work. We train our model on real-world clinical MRI raw data and evaluate it on two different tasks, skullstripping in k-Space and MRI reconstruction on the publicly available FastMRI knee dataset [13]. The code is made publicly available at <https://github.com/TIO-IKIM/PhaseGen>.

2 Material and Methods

2.1 (Complex Valued) Diffusion & Model Architecture

Complex valued neural networks (CVNNs) are a type of neural network that can process data that has both real and imaginary components, such as signals in the frequency domain. In CVNNs, the weights and activations are represented as complex numbers, allowing them to capture the phase information inherent in the data. This is particularly useful in applications like signal processing, where phase information is crucial for accurate analysis and reconstruction. The goal of diffusion models differs from other neural networks in that the network learns the diffusion process, introduced by gradually increasing noise to an input image. Given a data point x_0 , in the *forward diffusion process* during training, small amounts of Gaussian noise is added to the image in T consecutive time steps. The resulting data, superimposed with noise, x_t can be described with

$$q(x_t|x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t} \cdot x_{t-1}, \beta_t I), \quad (1)$$

where q is the forward diffusion process, $\mathcal{N}$ a normal distribution, I the identity matrix and β_t the scheduler, which defines the amount of noise ϵ added, commonly ranging from 0 to 1. During training, the model will then learn to predict the added noise in each timestep during the *reverse diffusion process* of the diffusion model parametrized by θ

$$p_\theta(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \Sigma_\theta(x_t, t)), \quad (2)$$

with μ_θ being the mean $\sqrt{1 - \beta_t} \cdot x_{t-1}$. To update the weights of the model, the loss is calculated for each timestep. During inference, the diffusion model then takes the normal noise distribution as input and via reverse diffusion samples a new data point.

Because MRI raw data is complex-valued, we have to adapt the diffusion and sampling process. With the magnitude already being present in the data, we only want to generate a corresponding phase. To simulate the nature of complex values z we define a complex-valued noise distribution $\mathcal{N}_z$. To focus on the phase to be generated, we adapt θ to have a magnitude of one, while the phase values are in the uniform distribution $\mathcal{U}$ with mean 0 and circular variance 1:

$$\mathcal{N}_z = e^{i\epsilon}; \quad \epsilon \sim \mathcal{U}(-\pi, \pi). \quad (3)$$

The adjusted complex-valued forward diffusion step can then be calculated as

$$z_t = (|z_{t-1}| \sqrt{\alpha_t} + |\epsilon| \sqrt{1 - \alpha_t}) \cdot \exp(\angle z_{t-1} + \angle \epsilon \sqrt{1 - \alpha_t}) \quad (4)$$

with $\alpha_t := 1 - \beta_t$, $|\cdot|$ denoting the magnitude and $\angle$ the phase of a complex number.

As underlying model architecture we use a complex-valued residual U-Net structure. The model takes as input during training z_t , to which noise is gradually added and in a second channel the magnitude of the input, which is supposed to be consistent throughout the process. During inference, the input consists of randomly sampled noise $\mathcal{N}_z$ and the magnitude in a second channel. The second channel is kept constant in every sampling step. The final output of PhaseGen is a complex-valued image, which can be transformed into the k-Space via Fourier transformation.

To evaluate whether magnitude-conditioned synthetic phase supports downstream k-Space learning, we will conduct multiple experiments. The first downstream task is skullstripping directly in the k-Space. This segmentation task, which separates the brain tissue from the skull in an image (in this case directly in the frequency domain), has already been presented in [14]. The accuracy is tested on a raw dataset gathered in-house. The comparison will be conducted on the model trained with artificial phase information created with the presented approach and other methods for data generation, such as random sampling. In the second experiment, we will use the artificial raw data to train a reconstruction algorithm on the publicly available FastMRI knee single coil dataset.

2.2 Datasets & Training

For the training of our diffusion model, we use an internal raw dataset gathered under IRB approval [24-11872-BO] on two separate MRI machines (Siemens 1.5T and 3T). In total we use 12071 2D raw data scans from 390 patients, consisting of different T_1 and T_2 sequences, as well as different resolutions. Each slice comprises 256×256 pixels. The datasets includes the following imaging sequences: TSE, TSE SPAIR, FL2D, TIRM and FLAIR with slice thickness between 2 and 8mm. For the first validation experiment we use an internal dataset, containing 21822 2D brain images from 150 patients, scanned with a 1.5 T and 3 T MRI machine, to train the skullstripping model. Some of these scans contain pathologies. Corresponding brain masks are already available as ground truth. Because this dataset only consists of magnitude image domain scans, we generate the phase data with our presented method and transform the data into the k-Space via the Fourier transformation. To investigate the benefit of training with artificially generated MRI raw data, we validate the trained model on real-world raw data, in total 14 volumes of individual patients, scanned on the two MRI machines mentioned above. The ground truth in the image domain is generated with the STAPLE algorithm [15], combining the results of three different skullstripping algorithms HD-BET [16], Synthstrip [17] and the de-identification tool presented in [18], on the image domain magnitude data. In the second validation experiment, performing the MRI reconstruction, we use the FastMRI single coil knee dataset [8] for training and testing.

For training, the Adam optimizer is used with an initial learning rate of $1e-4$. The learning rate is reduced exponentially with a gamma factor of 0.995, a beta

coefficient of 0.99 and an epsilon value of 1e-08. After each encoder convolution a dropout of 20% is chosen. The model is trained for 200 epochs with a batch size of 128. The overall training time takes 10 hours on an NVIDIA A100 GPU with 80GB of graphics memory. The diffusion model uses a cosine noise scheduler with 1000 timesteps and an ϵ of 0.008. The model has a total of 30.4 million parameters.

Qualitative examples of the generated phase can be found in the published code repository.

2.3 Skullstripping

In [14] the feasibility of performing skullstripping directly in the k-Space was shown, preserving valuable phase information for further downstream tasks. While the model performed well on synthetic datasets, generalizing to real-world clinical data proved challenging. In this first experiment we show the benefits of generating artificial phase for magnitude training data, leading to superior generalization on real-world raw data. The same model is trained on the magnitude image domain data and the artificial phase data, generated with the proposed complex-valued diffusion model. The model is then tested on real-world clinical data with original-phase values, to validate the benefit of training with the artificial phase data. The following data generation methods are compared:

- No phase data
- Naive phase generation
- Phase generation with the proposed diffusion model

The naive phase data generation, creates a synthetic phase by superimposing sinusoidal functions along both spatial dimensions and modulating it with the normalized magnitude to maintain anatomical structure correlation. Small random variations ($\sigma = 0.05$) were added to introduce realistic phase noise. The resulting synthetic phase can be described as

$$\phi(x, y) = \left[\sin\left(\frac{2\pi x}{N}\right) + \cos\left(\frac{2\pi y}{N}\right) \right] \cdot \hat{M}(x, y) + \eta(x, y), \quad (5)$$

where N is the image size, $\hat{M}$ the normalized magnitude and η the added noise. For each method, the resulting data is transformed into the k-Space. For all three methods, an extensive grid search was performed to find the best hyperparameters. The model architecture, a complex-valued residual U-Net, as presented in [14], is the same for all three methods. The results are shown in Fig. 2. The proposed method outperforms the other methods in both metrics, Dice similarity coefficient (DSC) and Hausdorff distance (HD). The model trained with the proposed method achieves a DSC of 80.1% and a HD of 1.534 pixel. The model trained with naive phase data achieves a DSC of 41.1% and a HD of 1.634 pixel. The model trained without phase data achieves a DSC of 40.1% and a HD of 1.577 pixel. It has to be noted, that the large difference in DSC is based on the fact, that both the model trained with no phase and the model with naive phase generation, are not able to generalize on the real-world data.




Fig. 2. Comparison of different phase data generation methods on the task of skull-stripping directly in the frequency domain. The model trained with the data generated by the proposed methods outperforms the other methods in both metrics, DSC and HD, showing superior segmentation performance.

2.4 MRI Reconstruction

The second validation experiment is the MRI reconstruction task. To reduce MRI acquisition times, the k-Space is often undersampled. This means that only a small part of the k-Space is acquired, while the rest is filled with zeros. This leads to a loss of information and thus to a lower image quality. The goal of this task is to reconstruct the missing k-Space lines, using the available data [19].

The FastMRI singlecoil knee dataset is used for training and testing a complex-valued residual U-Net, based on the architecture described above. We extend the existing model with data-consistency layers in the downsample path. We compare the model results of the model trained with the original raw data with the model trained on artificially generated phase data with the proposed method. The results are again compared with a model trained on naive phase data, generated as explained above. All experiments are conducted with a small model with roughly 209000 parameters (209k model) and a larger model with 3.3 million parameters (3M model). Following the official FastMRI challenge implementations, we use Cartesian undersampling masks with an 8% fully sampled center region for an undersampling rate of four and a 4% region for an undersampling rate of eight. To evaluate the performance of the reconstruction we use the commonly used metrics Peak Signal to Noise Ratio (PSNR) and Structural Similarity Index (SSIM), as well as the Mean Squared Error (MSE) and the Normalized Root Mean Squared Error (NRMSE). The results are shown in Tab.1 and Fig.3. Across both model sizes and undersampling factors, PhaseGen consistently outperforms naive phase generation, confirming that learning a struc-




Fig. 3. Results of the reconstruction task. Shown are the SSIM, PSNR, MSE and NRMSE for the model trained with original-phase data, naive phase generation and phase generation with the proposed method, as well as zerofilling as the baseline method. The results are shown for an undersampling factor of four. Higher values in the top row are better, while lower values in the bottom row are better.

tured, magnitude-conditioned phase distribution is more beneficial than heuristic synthesis. At $4\times$ undersampling, the performance gap to original-phase training narrows substantially with increasing model capacity, suggesting that PhaseGen-generated data captures phase statistics relevant to reconstruction. At $8\times$ undersampling, the advantage of PhaseGen over the naive baseline is maintained, though all synthetic-data models show a larger gap to the original-phase upper bound, reflecting the increased demand for accurate phase encoding at higher acceleration rates. To investigate the feasibility of combining PhaseGen-generated data with limited real phase data, the model is trained with varying proportions of original-phase data (1%, 10%, 15%, and 20%). This experiment is conducted with an undersampling factor of four. The results are shown in Tab.2. The 3M model trained with 15% - 20% of original-phase data achieves comparable results to the model trained with 100% of original-phase data. These results show that the proposed method can significantly reduce the amount of original phase data required for training, while achieving comparable results to a model trained only on original-phase data.

3 Discussion & Conclusion

We present a magnitude-conditioned generative phase modeling framework that enables synthesis of complex-valued MRI k-space data from magnitude-only

Table 1. Results of the reconstruction task. Compared is the same model with different sources of training data: naive phase generation, phase generation with PhaseGen and original-phase data, with zerofilling as the baseline. The numbers inside the bracket indicate the number of parameters of the model. The results are shown for an under-sampling factor of four and eight.

Training Data	SSIM (%) $\uparrow$	PSNR (dB) $\uparrow$	MSE $\downarrow$	NRMSE $\downarrow$
Undersampling x4				
<i>Original [209k]</i>	<i>69.10 $\pm$ 11.44</i>	<i>27.19 $\pm$ 2.22</i>	<i>0.002 $\pm$ 0.001</i>	<i>0.207 $\pm$ 0.061</i>
<i>Original [3.3M]</i>	<i>71.28 $\pm$ 11.76</i>	<i>27.98 $\pm$ 2.58</i>	<i>0.002 $\pm$ 0.001</i>	<i>0.190 $\pm$ 0.066</i>
Zerofilling	33.88 $\pm$ 7.89	18.62 $\pm$ 3.32	0.017 $\pm$ 0.011	0.553 $\pm$ 0.153
Naive [209k]	54.32 $\pm$ 10.64	18.31 $\pm$ 2.53	0.018 $\pm$ 0.011	0.606 $\pm$ 0.289
Naive [3.3M]	56.31 $\pm$ 14.44	21.81 $\pm$ 3.01	0.008 $\pm$ 0.007	0.403 $\pm$ 0.192
PhaseGen [209k]	59.73 $\pm$ 9.93	22.76 $\pm$ 2.94	0.007 $\pm$ 0.007	0.351 $\pm$ 0.134
PhaseGen [3.3M]	63.16 $\pm$ 10.87	23.95 $\pm$ 2.91	0.005 $\pm$ 0.004	0.301 $\pm$ 0.098
Undersampling x8				
<i>Original [209k]</i>	<i>64.74 $\pm$ 11.38</i>	<i>25.58 $\pm$ 2.46</i>	<i>0.003 $\pm$ 0.002</i>	<i>0.247 $\pm$ 0.069</i>
<i>Original [3.3M]</i>	<i>67.01 $\pm$ 11.97</i>	<i>26.26 $\pm$ 2.61</i>	<i>0.003 $\pm$ 0.002</i>	<i>0.231 $\pm$ 0.071</i>
Zerofilling	31.15 $\pm$ 8.47	17.12 $\pm$ 3.58	0.025 $\pm$ 0.016	0.673 $\pm$ 0.248
Naive [209k]	45.74 $\pm$ 11.11	20.41 $\pm$ 3.26	0.012 $\pm$ 0.008	0.465 $\pm$ 0.199
Naive [3.3M]	47.52 $\pm$ 13.41	20.32 $\pm$ 2.57	0.011 $\pm$ 0.007	0.454 $\pm$ 0.139
PhaseGen [209k]	54.57 $\pm$ 11.53	20.61 $\pm$ 3.28	0.011 $\pm$ 0.009	0.434 $\pm$ 0.096
PhaseGen [3.3M]	55.39 $\pm$ 12.11	21.22 $\pm$ 3.22	0.010 $\pm$ 0.010	0.405 $\pm$ 0.096

Table 2. Results of the reconstruction task. Compared are different amounts of original-phase data. The results are shown for an undersampling factor of four.

% of original data	SSIM (%) $\uparrow$	PSNR (dB) $\uparrow$	MSE $\downarrow$	NRMSE $\downarrow$
209k parameters				
1%	63.98 $\pm$ 10.37	24.26 $\pm$ 2.72	0.005 $\pm$ 0.004	0.301 $\pm$ 0.143
10%	67.42 $\pm$ 11.36	26.55 $\pm$ 2.47	0.003 $\pm$ 0.002	0.223 $\pm$ 0.071
15%	67.31 $\pm$ 11.07	26.22 $\pm$ 2.13	0.003 $\pm$ 0.001	0.231 $\pm$ 0.073
20%	67.62 $\pm$ 11.12	26.46 $\pm$ 2.61	0.003 $\pm$ 0.002	0.223 $\pm$ 0.058
100%	69.10 $\pm$ 11.44	27.19 $\pm$ 2.22	0.002 $\pm$ 0.001	0.207 $\pm$ 0.061
3.3M parameters				
1%	67.80 $\pm$ 10.78	25.78 $\pm$ 2.67	0.003 $\pm$ 0.003	0.249 $\pm$ 0.100
10%	69.65 $\pm$ 11.66	26.90 $\pm$ 2.44	0.002 $\pm$ 0.002	0.218 $\pm$ 0.082
15%	69.78 $\pm$ 11.67	27.22 $\pm$ 2.43	0.002 $\pm$ 0.002	0.207 $\pm$ 0.070
20%	71.26 $\pm$ 12.06	27.72 $\pm$ 2.65	0.002 $\pm$ 0.001	0.198 $\pm$ 0.071
100%	71.28 $\pm$ 11.76	27.98 $\pm$ 2.58	0.002 $\pm$ 0.001	0.190 $\pm$ 0.066

datasets. Our results show that structured phase modeling conditioned on magnitude supports effective downstream k-Space learning when original raw phase data is limited. For MRI reconstruction, training with a mixture of PhaseGen synthetic data and only 15–20% real data achieved performance comparable to training entirely on original-phase data, outperforming naive generation baselines. During our experiments we observed failure modes in form of phase artifacts at high-contrast tissue boundaries and reduced accuracy for sequences

outside the training distribution. While physics-informed approaches such as PISF [20] offer an alternative path to synthetic data generation for MRI reconstruction by encoding explicit acquisition priors, PhaseGen instead learns the phase distribution directly from data, requiring no sequence-specific physical modeling and thus generalizing more readily across protocols. Despite these promising results, the current model is limited to single-coil data, and with an inference time of roughly 10 seconds per slice, generating large datasets remains computationally expensive. Future work will therefore focus on optimizing inference speed and extending the framework to the multi-coil data commonly used in clinical settings. Our contribution lies in establishing the feasibility and practical value of magnitude-conditioned phase modeling for k-Space pretraining. Overall, magnitude-conditioned generative phase modeling provides a practical pathway to leverage magnitude-only MRI datasets for complex-valued learning, enabling data-efficient pretraining for raw-data-based downstream tasks.

Acknowledgments. This work received funding from the the Bruno & Helene Jöster Foundation, KITE (Plattform für KI-Translation Essen) from the REACT-EU initiative (<https://kite.ikim.nrw/>) and the Cancer Research Center Cologne Essen (CCCE).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Wojack, P., Fritz, J., Khodarahmi, I.: A Decade of Advancements in Musculoskeletal Imaging. *Investigative Radiology* 10–1097 (2024)
2. Egger, J., Gsaxner, C., Pepe, A., Pomykala, K.L., Jonske, F., Kurz, M., Li, J., Kleesiek, J.: Medical deep learning—A systematic meta-review. *Computer Methods and Programs in Biomedicine* 221, 106874 (2022)
3. Wang, J., Wang, S., Zhang, Y.: Deep learning on medical image analysis. *CAAI Transactions on Intelligence Technology* 10(1), 1–35 (2025)
4. Albuquerque, C., Henriques, R., Castelli, M.: Deep learning-based object detection algorithms in medical imaging: Systematic review. *Heliyon* 11(1) (2025)
5. Rempe, M., Hörst, F., Seibold, C., Hadaschik, B., Schlimbach, M., Egger, J., Kröninger, K., Breuer, F., Blaimer, M., Kleesiek, J.: Tumor Likelihood Estimation on MRI Prostate Data by Utilizing K-Space Information. *arXiv preprint arXiv:2407.06165* (2024)
6. Li, R., Pan, J., Zhu, Y., Ni, J., Rueckert, D.: Classification, Regression and Segmentation Directly from K-Space in Cardiac MRI. In: Xu, X., Cui, Z., Rekik, I., Ouyang, X., Sun, K. (eds.) *Machine Learning in Medical Imaging*, vol. 15241, pp. 31–41. Springer Nature Switzerland, Cham (2025)
7. Dishner, K.A., McRae-Posani, B., Bhowmik, A., Jochelson, M.S., Holodny, A., Pinker, K., Eskreis-Winkler, S., Stember, J.N.: A Survey of Publicly Available MRI Datasets for Potential Use in Artificial Intelligence Research. *Journal of Magnetic Resonance Imaging* 59(2), 450–480 (2024)
8. Zbontar, J., Knoll, F., Sriram, A., Murrell, T., Huang, Z., Muckley, M.J., Defazio, A., Stern, R., Johnson, P., Bruno, M., Parente, M., Geras, K.J., Katsnelson, J., Chandarana, H., Zhang, Z., Drozdal, M., Romero, A., Rabbat, M., Vincent, P.,

- Yakubova, N., Pinkerton, J., Wang, D., Owens, E., Zitnick, C.L., Recht, M.P., Sodickson, D.K., Lui, Y.W.: fastMRI: An Open Dataset and Benchmarks for Accelerated MRI. arXiv preprint arXiv:1811.08839 (2018)
9. Tibrewala, R., Dutt, T., Tong, A., Ginocchio, L., Lattanzi, R., Keerthivasan, M.B., Baete, S.H., Chopra, S., Lui, Y.W., Sodickson, D.K., et al.: FastMRI Prostate: A public, biparametric MRI dataset to advance machine learning for prostate cancer imaging. *Scientific Data* 11(1), 404 (2024)
 10. Solomon, E., Johnson, P.M., Tan, Z., Tibrewala, R., Lui, Y.W., Knoll, F., Moy, L., Kim, S.G., Heacock, L.: fastMRI Breast: A publicly available radial k-space dataset of breast dynamic contrast-enhanced MRI. *Radiology: Artificial Intelligence* 7(1), e240345 (2025)
 11. Singh, D., Monga, A., De Moura, H.L., Zhang, X., Zibetti, M.V.W., Regatte, R.R.: Emerging Trends in Fast MRI Using Deep-Learning Reconstruction on Undersampled k-Space Data: A Systematic Review. *Bioengineering* 10(9), 1012 (2023)
 12. Croitoru, F.-A., Hondru, V., Ionescu, R.T., Shah, M.: Diffusion Models in Vision: A Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45(9), 10850–10869 (2023)
 13. Knoll, F., Zbontar, J., Sriram, A., Muckley, M.J., Bruno, M., Defazio, A., Parente, M., Geras, K.J., Katsnelson, J., Chandarana, H., Zhang, Z., Drozdal, M., Romero, A., Rabbat, M., Vincent, P., Pinkerton, J., Wang, D., Yakubova, N., Owens, E., Zitnick, C.L., Recht, M.P., Sodickson, D.K., Lui, Y.W.: fastMRI: A Publicly Available Raw k-Space and DICOM Dataset of Knee Images for Accelerated MR Image Reconstruction Using Machine Learning. *Radiology: Artificial Intelligence* 2(1), e190007 (2020)
 14. Rempe, M., Mentzel, F., Pomykala, K.L., Haubold, J., Nensa, F., Kroeninger, K., Egger, J., Kleesiek, J.: K-Strip: A Novel Segmentation Algorithm in k-Space for the Application of Skull Stripping. *Computer Methods and Programs in Biomedicine* 243, 107912 (2024)
 15. Warfield, S.K., Zou, K.H., Wells, W.M.: Simultaneous Truth and Performance Level Estimation (STAPLE): An Algorithm for the Validation of Image Segmentation. *IEEE Transactions on Medical Imaging* 23(7), 903–921 (2004)
 16. Isensee, F., Schell, M., Pflueger, I., Brugnara, G., Bonekamp, D., Neuberger, U., Wick, A., Schlemmer, H.-P., Heiland, S., Wick, W., Bendszus, M., Maier-Hein, K.H., Kieringeder, P.: Automated Brain Extraction of Multisequence MRI Using Artificial Neural Networks. *Human Brain Mapping* 40(17), 4952–4964 (2019)
 17. Hoopes, A., Mora, J.S., Dalca, A.V., Fischl, B., Hoffmann, M.: SynthStrip: Skull-Stripping for Any Brain Image. *NeuroImage* 260, 119474 (2022)
 18. Rempe, M., Heine, L., Seibold, C., Hö rst, F., Kleesiek, J.: De-identification of medical imaging data: a comprehensive tool for ensuring patient privacy. *European Radiology* 1–10 (2025)
 19. Hyun, C. M., Kim, H. P., Lee, S. M., Lee, S., Seo, J. K.: Deep learning for under-sampled MRI reconstruction. *Physics in Medicine & Biology*, 63(13) (2018)
 20. Wang, Z., Yu, X., Wang, C., Chen, W., Wang, J., Chu, Y. H., ... & Qu, X. One for multiple: Physics-informed synthetic data boosts generalizable deep learning for fast MRI reconstruction. *Medical Image Analysis*, 103, 103616. (2025)