

Gabor Primitives for Accelerated Cardiac Cine MRI Reconstruction

Wenqi Huang^{1,2}✉, Veronika Spieker^{2,3,4}, Nil Stolt-Ansó^{1,2,5}, Natascha Niessen^{2,6}, Maik Dannecker^{1,2}, Sevgi Gokce Kafali^{1,2}, Sila Kurugol⁴, Julia A. Schnabel^{2,3,5,7}, and Daniel Rueckert^{1,2,5,8}

¹ Chair for AI in Healthcare and Medicine, Technical University of Munich (TUM) and TUM University Hospital, Munich, Germany

² School of Computation and Information Technology, Technical University of Munich, Munich, Germany

³ Institute of Machine Learning in Biomedical Imaging, Helmholtz Munich, Neuherberg, Germany

⁴ Department of Radiology, Boston Children's Hospital and Harvard Medical School, Boston, USA

⁵ Munich Center for Machine Learning, Technical University of Munich, Munich, Germany

⁶ GE HealthCare, Munich, Germany

⁷ School of Biomedical Engineering and Imaging Sciences, King's College London, London, United Kingdom

⁸ Department of Computing, Imperial College London, London, United Kingdom
wenqi.huang@tum.de

Abstract. Accelerated cardiac cine MRI requires reconstructing spatiotemporal images from highly undersampled k -space data. Implicit neural representations (INRs) enable scan-specific reconstruction without large training datasets, but encode content implicitly in network weights without physically interpretable parameters. Gaussian primitives provide an explicit and geometrically interpretable alternative, but their spectra are confined near the k -space origin, limiting high-frequency representation. We propose Gabor primitives for MRI reconstruction, modulating each Gaussian envelope with a complex exponential to place its spectral support at an arbitrary k -space location, enabling efficient representation of both smooth structures and sharp boundaries. To exploit spatiotemporal redundancy in cardiac cine, we decompose per-primitive temporal variation into a low-rank geometry basis capturing cardiac motion and a signal-intensity basis modeling contrast changes. Experiments on cardiac cine data with Cartesian and radial trajectories show that Gabor primitives consistently outperform compressed sensing, Gaussian primitives, and hash-grid INR baselines, while providing a compact, continuous-resolution representation with physically meaningful parameters.⁹

Keywords: MRI Reconstruction · Gaussian Splatting · Gabor Primitives · Cardiac Cine MRI · Implicit Neural Representations

⁹ Code available at <https://github.com/wenqihuang/FastGaborMRI>.

1 Introduction

Cardiac cine MRI requires high spatial and temporal resolution, yet the achievable resolution is constrained by scan duration, which is limited by MR physics and patients’ tolerance [18, 23]. Undersampling k -space is therefore essential, but results in an ill-posed inverse reconstruction problem that requires effective priors or regularization.

Compressed sensing (CS) and low-rank methods exploit sparsity and temporal redundancy [7, 16], but their hand-crafted regularizers may struggle with complex, patient-specific motion and fine anatomical detail. Supervised deep learning methods can learn stronger data-driven priors [9, 17], yet require large training datasets and risk hallucination on out-of-distribution cases [2, 11]. Implicit neural representations (INRs) offer a scan-specific alternative requiring no external data [6, 10, 20], but encode information implicitly in network weights with no physical interpretability.

In computer vision, 3D Gaussian splatting (3DGS) [12] represents scenes as collections of Gaussian primitives with learnable position, shape, color and opacity. This idea has recently been adapted to MRI reconstruction [19, 21]. However, individual Gaussian primitives are inherently smooth in image space, making sharp features such as tissue boundaries costly to represent. In the frequency domain, their spectra remain centered at the k -space origin, as spatial translation introduces only a phase shift but not a frequency shift. Computer vision has addressed this high-frequency limitation using frequency-modulated primitives such as wavelet [27] and Gabor splatting [24, 25, 28], but these are designed for non-negative real-valued signals and do not directly extend to the complex-valued MRI setting.

In this work, we instead modulate each Gaussian with a complex exponential, yielding a freely positionable spectral component per primitive that can capture high-frequency content directly rather than through superposition of many narrow Gaussians. The formulation generalizes standard Gaussian primitives, which are recovered when the modulation frequency is zero, so low-frequency representation is preserved. We further propose a two-component low-rank temporal model that separates geometry dynamics from signal-intensity variations, and validate the framework on cardiac cine MRI with both Cartesian and non-Cartesian trajectories. Our contributions are:

1. We introduce a novel complex-valued Gabor primitive formulation for MRI reconstruction, where each primitive carries a freely positionable k -space spectral component. This generalizes standard Gaussian primitives and enables efficient representation of both smooth anatomy and sharp boundaries.
2. We design a two-component low-rank temporal model that splits per-primitive dynamics into geometry and signal-intensity bases, providing a structured and compact parametrization of cardiac motion and contrast changes.
3. Across multiple cardiac cine datasets with Cartesian and radial trajectories at high acceleration factors, we demonstrate consistent improvements over compressed sensing, Gaussian primitive, and hash-grid INR baselines in all evaluation settings.




Fig. 1. Left: Gaussian vs. Gabor primitives in image space and k -space. A Gaussian’s spectral support is fixed at the k -space origin; a Gabor primitive shifts it to ξ_i via complex-exponential modulation, enabling more efficient frequency coverage. **Right:** Cardiac cine image is modeled as a mixture of time-varying Gabor primitives. Geometry parameters (μ, s, θ, ξ) and complex weights w are generated from low-rank geometry (blue) and intensity (orange) bases. Primitives are rasterized, passed through a multi-coil forward model, and fitted to acquired k -space data end-to-end.

2 Methods

2.1 Gabor Primitive Basis

A standard Gaussian primitive centered at μ_n with covariance Σ_n defines a smooth spatial basis function. Its Fourier transform has a Gaussian envelope centered at the k -space origin regardless of the primitive’s spatial position, so representing high-frequency content requires superposition of many narrow primitives. To overcome this limitation, we modulate each Gaussian with a complex exponential to form a Gabor primitive [4, 8, 25]:

$$P_n(\mathbf{r}) = \underbrace{\exp(i 2\pi \xi_n \cdot (\mathbf{r} - \mu_n))}_{\text{complex exponential}} \cdot \underbrace{\exp\left(-\frac{1}{2}(\mathbf{r} - \mu_n)^\top \Sigma_n^{-1}(\mathbf{r} - \mu_n)\right)}_{\text{Gaussian envelope}}, \quad (1)$$

where $\mathbf{r} \in \mathbb{R}^2$ is the spatial coordinate, $\mu_n \in \mathbb{R}^2$ is the center position, $\xi_n \in \mathbb{R}^2$ is the modulation frequency, and $\Sigma_n = R_{\theta_n} \text{diag}(s_{n,1}^2, s_{n,2}^2) R_{\theta_n}^\top$ is the covariance matrix parameterized by rotation angle θ_n and two scale parameters $s_{n,k} > 0$ along each principal axis. The Fourier transform of P_n is:

$$\hat{P}_n(\mathbf{k}) = \frac{2\pi}{\sqrt{|\Sigma_n^{-1}|}} \exp\left(-2\pi^2(\mathbf{k} - \xi_n)^\top \Sigma_n(\mathbf{k} - \xi_n)\right) e^{-i2\pi \mathbf{k} \cdot \mu_n}, \quad (2)$$

where $\mathbf{k} \in \mathbb{R}^2$ is the k -space coordinate. Thus, each primitive produces a single Gaussian blob in k -space centered at the modulation frequency ξ_n rather than

at the origin (Fig. 1, left), reducing to a standard Gaussian when $\boldsymbol{\xi}_n = \mathbf{0}$. By the Fourier scaling property, the k -space blob inherits the same orientation but with reciprocal widths. Compared to standard Gaussians whose spectra are all anchored at the k -space origin, Gabor primitives place spectral support at arbitrary k -space locations, reducing spectral overlap and improving representation efficiency.

2.2 Spatiotemporal Forward Model

We model the cardiac cine MR image across T frames as a sum of N Gabor primitives whose weights and geometry may vary across frames:

$$x_t(\mathbf{r}) = \sum_{n=1}^N w_{n,t} P_n(\mathbf{r}; \boldsymbol{\mu}_{n,t}, \mathbf{s}_{n,t}, \theta_{n,t}, \boldsymbol{\xi}_{n,t}), \quad t = 1, \dots, T, \quad (3)$$

where $w_{n,t} \in \mathbb{C}$ is the complex weight. We model this temporal variation through two low-rank temporal bases (Fig. 1, right): a *geometry basis* capturing cardiac motion and a *intensity basis* modeling signal-intensity variations.

Geometry Basis. Each geometric parameter is decomposed into a static component plus a low-rank perturbation from a shared basis $\mathbf{V}_g \in \mathbb{R}^{T \times R_g}$:

$$\begin{aligned} \boldsymbol{\mu}_{n,t} &= \boldsymbol{\mu}_n + \mathbf{C}_n^{[\mu]} \mathbf{v}_{g,t}, & \theta_{n,t} &= \theta_n + \mathbf{c}_n^{[\theta]} \cdot \mathbf{v}_{g,t}, \\ \log \mathbf{s}_{n,t} &= \log \mathbf{s}_n + \mathbf{C}_n^{[s]} \mathbf{v}_{g,t}, & \boldsymbol{\xi}_{n,t} &= \boldsymbol{\xi}_n + \mathbf{C}_n^{[\xi]} \mathbf{v}_{g,t}, \end{aligned} \quad (4)$$

where $\mathbf{v}_{g,t} \in \mathbb{R}^{R_g}$ is the t -th row of $\mathbf{V}_g$, $\log$ denotes the element-wise logarithm, and the $\mathbf{C}_n^{[\cdot]}$ are per-primitive coefficient matrices that determine how each primitive couples to the shared motion basis.

Contrast Basis. In cardiac cine MRI, through-plane motion causes anatomical structures to move in and out of the imaging plane, producing frame-to-frame intensity variations beyond what in-plane geometry changes can explain. The weight $w_{n,t}$ captures these variations:

$$w_{n,t} = \mathbf{u}_n \cdot \mathbf{v}_{c,t} + \mathbf{c}_n^{[w]} \cdot \mathbf{v}_{g,t}, \quad (5)$$

where $\mathbf{v}_{c,t} \in \mathbb{C}^{R_c}$ is the t -th row of an independent intensity basis $\mathbf{V}_c \in \mathbb{C}^{T \times R_c}$, $\mathbf{u}_n \in \mathbb{C}^{R_c}$ are per-primitive intensity coefficients, and $\mathbf{c}_n^{[w]} \in \mathbb{C}^{R_g}$ couples the weight to the geometry basis to model motion-induced intensity changes such as partial-volume and flow effects. The resulting weight matrix $\mathbf{W} \in \mathbb{C}^{N \times T}$ has rank at most $R_c + R_g$, imposing a low-rank temporal structure directly in primitive parameter space.

Multi-Coil Forward Model. Multi-coil k -space data can be derived from the predicted images x_t by applying the forward model for coil j at frame t :

$$\hat{y}_{j,t} = \mathcal{F}_\Omega\{\mathcal{S}_j x_t\}, \quad (6)$$

where $\mathcal{S}_j$ denotes a pre-estimated coil sensitivity map and $\mathcal{F}_\Omega$ is the Fourier transform evaluated on the acquired trajectory Ω , implemented as a masked FFT for Cartesian sampling or a non-uniform FFT (NUFFT) for non-Cartesian trajectories. In practice, we first rasterize Eq. (3) on a Cartesian grid using the frame-specific geometry from Eq. (4) and Eq. (5), multiply the results by each coil sensitivity, and apply $\mathcal{F}_\Omega$ (Eq. (6)). Overall, the model is optimized using the composite loss:

$$\mathcal{L} = \sum_{j=1}^{N_c} \sum_{t=1}^T \|y_{j,t} - \hat{y}_{j,t}\|_2^2 + \lambda_s \sum_{n,t} |w_{n,t}| + \lambda_t \sum_{t=1}^{T-1} \|x_{t+1} - x_t\|_1, \quad (7)$$

combining data fidelity between the predicted and measured k -space, ℓ_1 sparsity on primitive weights $w_{n,t}$ (λ_s), and temporal total variation (λ_t). All learnable quantities, including static primitive parameters, coefficient matrices, and temporal bases, are optimized jointly end-to-end.

2.3 Training and Inference

Primitive centers μ_n are initialized on a uniform grid covering the imaging domain; modulation frequencies ξ_n are initialized near zero and evolve freely during optimization, subject to a fixed soft bound on $\|\xi_n\|$ that keeps modulation frequencies well below the rasterization Nyquist limit. Temporal coefficient matrices are initialized to zero, meaning that the model starts static and gradually discovers motion. We adopt adaptive density control [12], periodically pruning low-contribution primitives and splitting those with large gradients. Optimization uses Adam [13] with per-parameter-group learning rates and cosine annealing. At inference, the continuous representation can be evaluated on any spatial grid.

3 Experiments

Datasets. We evaluate on two publicly available cardiac cine MRI datasets, both released with prior ethics approval and de-identification by the original studies. **(1) Cartesian:** 99 fully sampled bSSFP cardiac cine acquisitions (55 at 1.5 T, 44 at 3 T, 208×256 , ≈ 19 frames, 15–18 coils) from [3], retrospectively undersampled at $R=12$ and $R=16$ with VISTA masks [1]. Coil sensitivities were estimated via ESPIRiT [22]. **(2) Radial:** 102 radial bSSFP acquisitions at 3 T from [5] (208×208 , 25 cardiac phases, coils compressed to 8), retrospectively subsampled to 14 spokes per frame ($R \approx 23$). Coil sensitivities were estimated by ESPIRiT applied to a low-resolution image obtained by gridding of the time-averaged spokes, and references were obtained via CG-SENSE from all spokes.

Table 1. Quantitative comparison on cardiac cine MRI ($N=99$ Cartesian, $N=102$ radial). PSNR in dB, Time in minutes. $\rho = N_{\text{params}}/(2HWT)$: ratio of learnable parameters to real-valued image degrees of freedom. Best in **bold**, second best underlined.

Method	$\rho \downarrow$	Cartesian $R=12$				Cartesian $R=16$				Radial $R \approx 23$			
		PSNR $\uparrow$	SSIM $\uparrow$	FSIM $\uparrow$	Time $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	FSIM $\uparrow$	Time $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	FSIM $\uparrow$	Time $\downarrow$
L+S [16]	—	37.15	0.940	0.850	<u>0.3</u>	32.19	0.824	0.777	<u>0.4</u>	35.82	0.922	0.810	<u>1.5</u>
PICS [14]	—	<u>42.21</u>	<u>0.972</u>	0.886	0.1	<u>41.05</u>	0.966	0.875	0.1	35.19	0.905	0.793	1.1
Hash-INR [15]	2.60	39.63	0.921	0.841	4.6	37.77	0.893	0.815	4.6	34.00	0.849	0.767	8.0
Gaussian	<u>0.42</u>	41.50	0.972	0.874	1.2	40.67	<u>0.967</u>	0.864	1.2	<u>36.67</u>	<u>0.942</u>	<u>0.810</u>	3.4
Gabor (ours)	0.41	42.61	0.974	<u>0.883</u>	2.5	41.39	0.968	<u>0.871</u>	2.5	37.53	0.947	0.830	4.2

Baselines. We compare against two conventional methods and two scan-specific learned representations: (1) **L+S** [16], robust PCA decomposition into low-rank plus sparse components; (2) **PICS** [14], iterative parallel imaging and compressed sensing with total variation regularization; (3) **Hash-INR**, a scan-specific implicit neural representation with multi-resolution hash-grid encoding [15]; and (4) **Gaussian primitives**, our framework with $\xi_n \equiv \mathbf{0}$, serving as an ablation baseline.

Implementation Details. All methods were implemented in PyTorch on an NVIDIA A6000 GPU, with custom CUDA kernels for primitive rasterization. To match total parameter counts, Gabor used $N=15,000$ initial primitives growing to 20,000, while Gaussian used 19,000 growing to 21,000. Both used temporal ranks $R_g=6$, $R_c=4$ with geometry-contrast coupling, adaptive density control (every 300 iterations, 10–60% of training), and 5,000 iterations ($\lambda_s=10^{-5}$, $\lambda_t=10^{-2}$). Hash-INR used 16 hash levels (2^{19} map size), a 2-hidden-layer MLP (128 nodes) and trained for 5,000 iterations. Quality was evaluated by peak signal-to-noise ratio (PSNR), structural similarity (SSIM), and feature similarity (FSIM) [26]. Temporal ranks and regularization weights were selected via grid search on a representative subset of five scans.

4 Results

Quantitative Comparison. Table 1 summarizes reconstruction quality across all test slices. Gabor primitives achieve the highest PSNR and SSIM across all three settings, with consistent gains over Gaussian (+1.11/+0.72/+0.86 dB). While PICS achieves marginally higher FSIM on Cartesian data, Gabor primitives outperform all competing methods on radial data ($R \approx 23$), including a +2.34 dB gain over PICS. This highlights the stronger implicit structural prior of Gabor compared to TV regularization under severe undersampling, while offering the most compact image representation ($\rho < 0.5$). Hash-INR requires $6 \times$ more parameters ($\rho=2.60$) yet ranks last among learned methods. Gabor primitives are slower than Gaussian because they capture high frequencies via modulation rather than spatial narrowing, resulting in spatially wider primitives on




Fig. 2. Reconstruction comparison on Cartesian (top) and radial (bottom) cardiac cine. Each block shows a representative frame with metrics, zoomed ROI with error map, and y - t profile with temporal error.

average. During rasterization, each pixel therefore overlaps with more Gabor than Gaussian primitives, increasing per-pixel computation.

Qualitative Comparison. Figure 2 compares representative reconstructions across all methods. Voxel-based methods (PICS, L+S, Hash-INR) show visible background noise in the error maps, since their per-voxel degrees of freedom allow fitting noise alongside signal. Primitive methods suppress this via locally smooth, compact representations ($\rho < 0.5$). L+S blurs temporal motion boundaries in the y - t profiles, and all non-Gabor methods degrade more on radial data. Gaussian shows elevated error at tissue boundaries under high acceleration (radial zoom), where the absence of carrier frequency ξ limits edge representation. Gabor yields the lowest spatial and temporal error across both settings.

Representation Properties. Beyond reconstruction quality, Gabor primitives provide unique capabilities over grid-based and Gaussian methods.

k -space coverage. Figure 3a shows that learned Gabor primitives distribute their center frequencies ξ_n across k -space to match the image’s spectral content; the overlaid 3σ support ellipses illustrate the anisotropic spectral footprint of individual primitives. In contrast, Gaussian primitives are anchored at the origin. This spectral advantage is confirmed quantitatively in Fig. 3b: across the radial




Fig. 3. Unique capabilities of Gabor primitives. **(a)** k -space primitive distribution: each dot is a primitive’s center frequency ξ_n with 3σ support ellipses. **(b)** Frequency-band PSNR (low, mid, high) on the radial dataset, averaged over all subjects. **(c)** Frequency decomposition: primitives partitioned by $|\xi_n|$ into low- ($< \frac{1}{4}$) and high-frequency ($\geq \frac{1}{4}$) groups and their sum. **(d)** $4\times$ super-resolution on the marked ROI; bicubic upsampling of the reference shown for comparison.

cohort, all methods achieve similar PSNR in the low-frequency band, while Gabor yields the largest gains in the mid- and high-frequency bands where Gaussian and L+S fall short.

Frequency decomposition. Partitioning primitives by $|\xi_n|$ yields a natural spectral decomposition (Fig. 3c): low-frequency primitives capture smooth anatomy, while high-frequency ones capture edges (no Gaussian counterpart, where $|\xi_n| = 0$).

Continuous-resolution evaluation. As a continuous representation, Gabor primitives can be evaluated on an arbitrarily dense spatial grid without re-training (Fig. 3d). Since no true high-resolution acquisition is available for these datasets, we show a $4\times$ visual comparison against a bicubic upsampling of the reference; the comparison is qualitative rather than a quantitative super-resolution claim. Gabor primitives recover sharper structures than Gaussian primitives, which lack the carrier frequency ξ needed to extrapolate high-frequency detail, while Hash-INR produces visible raster artifacts at the increased resolution because it does not inherently guarantee spatial continuity and overfits to the sampled grid.

5 Discussion and Conclusion

We proposed Gabor primitives for MRI reconstruction, where complex exponential modulation gives each primitive a freely positionable k -space spectral component, enabling efficient high-frequency coverage that Gaussian primitives cannot achieve. A two-component temporal model factorizes per-primitive dynamics into geometry and intensity bases, imposing a structured low-rank prior in primitive parameter space. Across multiple cardiac cine settings with Cartesian and radial trajectories, Gabor primitives consistently outperform all baselines. The explicit parameterization also offers physical interpretability, such as spectral decomposition by modulation frequency, and could provide compact descriptors for downstream tasks such as motion quantification in the future.

Several limitations remain. As a scan-specific method, each acquisition requires individual optimization (2–4 minutes per scan), roughly twice the Gaussian baseline because Gabor primitives have wider spatial support and overlap more pixels during rasterization despite tile-based culling; further accelerations such as hierarchical rasterization and mixed-precision arithmetic are directly applicable and not yet exploited. The formulation operates in 2D; extending to 3D and joint spatiotemporal frequency modulation could improve temporal modeling. Validation on additional anatomies and clinical diagnostic evaluation are needed for translation.

Acknowledgements

This work was partially funded by the European Research Council (ERC) under Grant Deep4MI (No. 884622). The authors acknowledge GE HealthCare, Munich, for providing data for early-stage testing in this study. V.S. was supported by the Add-On Fellowship of the Joachim Herz Foundation. S.G.K. has been sponsored by the Alexander von Humboldt Foundation. S.K. was supported by the National Institute of Diabetes and Digestive and Kidney Diseases (NIDDK) of the National Institutes of Health under Award No. R01DK125561 and by the Massachusetts AI Hub Award.

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

References

1. Ahmad, R., Xue, H., Giri, S., Ding, Y., Craft, J., Simonetti, O.P.: Variable density incoherent spatiotemporal acquisition (vista) for highly accelerated cardiac MRI. *Magnetic Resonance in Medicine* **74**(5), 1266–1278 (2015)
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. Chen, C., Liu, Y., Schniter, P., Tong, M., Zareba, K., Simonetti, O., Potter, L., Ahmad, R.: OCMR (v1.0)—open-access multi-coil k-space dataset for cardiovascular magnetic resonance imaging. arXiv preprint arXiv:2008.03410 (2020)
4. Daugman, J.G.: Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters. *Journal of the Optical Society of America A* **2**(7), 1160–1169 (1985)
5. El-Rewaidy, H.: Replication Data for: Multi-Domain Convolutional Neural Network (MD-CNN) For Radial Reconstruction of Dynamic Cardiac MRI (2020)
6. Feng, J., Feng, R., Wu, Q., Shen, X., Chen, L., Li, X., Feng, L., Chen, J., Zhang, Z., Liu, C., Zhang, Y., Wei, H.: Spatiotemporal implicit neural representation for unsupervised dynamic MRI reconstruction. *IEEE Transactions on Medical Imaging* **44**(5), 2143–2156 (2025)
7. Feng, L.: Golden-angle radial MRI: basics, advances, and applications. *Journal of Magnetic Resonance Imaging* **56**(1), 45–62 (2022)
8. Gabor, D.: Theory of communication. *Journal of the Institution of Electrical Engineers—Part III: Radio and Communication Engineering* **93**(26), 429–457 (1946)
9. 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)
10. Huang, W., Li, H.B., Pan, J., Cruz, G., Rueckert, D., Hammernik, K.: Neural implicit k-space for binning-free non-cartesian cardiac MR imaging. In: *International Conference on Information Processing in Medical Imaging*. pp. 548–560. Springer (2023)
11. Johnson, P.M., Jeong, G., Hammernik, K., Schlemper, J., Qin, C., Duan, J., Rueckert, D., et al.: Evaluation of the robustness of learned MR image reconstruction to systematic deviations between training and test data for the models from the fastMRI challenge. In: *Machine Learning for Medical Image Reconstruction (MLMIR), MICCAI Workshop. Lecture Notes in Computer Science*, vol. 12964, pp. 25–34. Springer (2021)
12. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3D Gaussian Splatting for real-time radiance field rendering. *ACM Transactions on Graphics* **42**(4), 1–14 (2023)
13. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: *International Conference on Learning Representations (ICLR)* (2015)
14. Lustig, M., Donoho, D., Pauly, J.M.: Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic Resonance in Medicine* **58**(6), 1182–1195 (2007)
15. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (TOG)* **41**(4), 1–15 (2022)
16. Otazo, R., Candès, E., Sodickson, D.K.: Low-rank plus sparse matrix decomposition for accelerated dynamic MRI with separation of background and dynamic components. *Magnetic Resonance in Medicine* **73**(3), 1125–1136 (2015)
17. 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)
18. Rajiah, P.S., François, C.J., Leiner, T.: Cardiac MRI: State of the art. *Radiology* **307**(3), e223008 (2023)
19. Singh, I., Dupuis, A., Kukran, S., Badve, C., Griswold, M.: Back to basis: Spatially continuous MRI with adaptive Gaussians. In: *ISMRM Workshop on Data Sampling and Image Reconstruction. Sedona, AZ, USA* (2026)

20. Spieker, V., Huang, W., Eichhorn, H., Stelter, J., Weiss, K., Zimmer, V.A., Braren, R.F., Karampinos, D.C., Hammernik, K., Schnabel, J.A.: ICoNIK: Generating respiratory-resolved abdominal MR reconstructions using neural implicit representations in k-space. In: Deep Generative Models (DGM4MICCAI), MICCAI Workshop. pp. 183–192. Lecture Notes in Computer Science, Springer (2024)
21. Terpstra, M., van den Berg, C.: Fast undersampled dynamic MRI reconstruction using explicit representation learning with Gaussian splatting. In: ISMRM Workshop on Data Sampling and Image Reconstruction. Sedona, AZ, USA (2026)
22. Uecker, M., Lai, P., Murphy, M.J., Virtue, P., Elad, M., Pauly, J.M., Vasanawala, S.S., Lustig, M.: ESPIRiT - an eigenvalue approach to autocalibrating parallel MRI: where SENSE meets GRAPPA. *Magnetic Resonance in Medicine* **71**(3), 990–1001 (2014)
23. Wang, X., Uecker, M., Feng, L.: Fast real-time cardiac MRI: A review of current techniques and future directions. *Investigative Magnetic Resonance Imaging* **25**(4), 252–265 (2021)
24. Watanabe, H., Tojo, K., Umetani, N.: 3D gabor splatting: Reconstruction of high-frequency surface texture using gabor noise. In: Eurographics 2025 Short Papers (2025)
25. Wurster, S.W., Zhang, R., Zheng, C.: Gabor splatting for high-quality gigapixel image representations. In: ACM SIGGRAPH 2024 Posters (2024)
26. Zhang, L., Zhang, L., Mou, X., Zhang, D.: FSIM: A feature similarity index for image quality assessment. *IEEE Transactions on Image Processing* **20**(8), 2378–2386 (2011)
27. Zhang, W., Zhu, H., Wu, D., Kang, D., Bao, L., Cao, X., Ma, Z.: WIPES: Wavelet-based visual primitives. In: IEEE/CVF International Conference on Computer Vision (2025)
28. Zhou, J., Huang, Y., Dai, W., Zou, J., Zheng, Z., Kan, N., Li, C., Xiong, H.: 3DGabSplat: 3D gabor splatting for frequency-adaptive radiance field rendering. In: Proceedings of the 33rd ACM International Conference on Multimedia (2025)