

Implicit representations are dead. Long live explicit primitives!

Nil Stolt-Ansó ^{*1}, Maik Dannecker^{*2}, Wenqi Huang¹,
Andras Jakab², and Daniel Rueckert^{1,3}

¹ Technical University Munich, Germany

² University of Zurich, Switzerland

³ Imperial College London, United Kingdom
`nil.stolt@tum.de` – `m.dannecker@tum.de`

Abstract. Continuous parameterization of medical data has emerged as a powerful paradigm for resolution-independent image representation. While Implicit Neural Representations offer high fidelity and compact storage, their reliance on global Multi-Layer Perceptrons incurs sizeable computational costs, large memory requirements, and extensive optimization times. As medical imaging trends towards ever-more detailed, high-resolution volumes, these costs impose significant bottlenecks in the applicability of implicit approaches. Recently, explicit Gaussian-based primitives have revolutionized the representation learning paradigm by trading deep network evaluations for localized, rasterization-friendly primitives. In this paper, we present a comprehensive, cross-dimensional evaluation of Gaussian representations against implicit approaches for medical imaging applications. First, we outline a theoretical overview on the mathematical properties offered by explicit primitives beyond what is capable under the implicit neural paradigm. Subsequently, we benchmark the computational performance on two demanding image datasets: 2D microscopy histology and 3D lung Computed Tomography (CT). Our experiments demonstrate that Gaussian representations consistently match or surpass reconstruction metrics compared to implicit methods across all compression factors, while displaying significantly lower optimization times, and memory requirements. Together with the compelling mathematical properties offered by explicit primitives, these findings motivate the wider adoption of Gaussian representations and position them as an attractive direction for future research in medical imaging.

Keywords: Explicit Representations · Gaussian Primitives · Implicit Neural Representations · Medical Image Reconstruction.

1 Introduction

Medical images are acquired as discrete voxels of intensities, sampled along structures such as grids. While standardized, these formats place a limit in spatial

^{*} Equal contribution

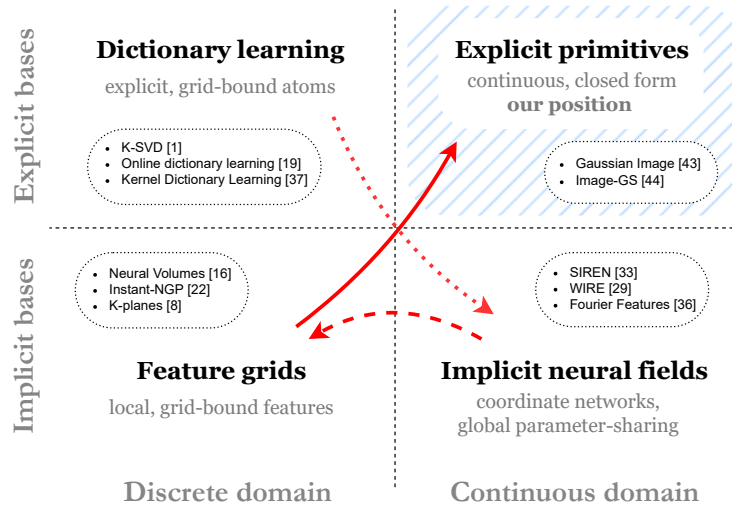


Fig. 1. Representation paradigms positioned by domain (discrete $\rightarrow$ continuous) and basis type (implicit $\rightarrow$ explicit). Red arrows indicate the direction of research in recent years: (1-dotted) movement from explicit dictionaries towards neural representations, (2-dashed) shift towards feature-grids, (3-solid) switch to explicit primitives.

and temporal resolution in that a scan’s intrinsic detail is directly coupled to its memory footprint. Off-grid methods instead reframe data as continuous functions, being mostly unaffected by sampling structure and enabling arbitrary upsampling. Crucially, how this function is parameterized dictates optimization speed, memory dynamics, and ultimately clinical viability.

Traditionally, dictionary learning and sparse coding methods [1, 19, 37] represented images as explicit bases, with grid-bound weights dictating the components of the signal at a given location. With the popularity of deep learning, Implicit Neural Representations (INRs) freed the basis from the grid by storing the signal in the weights of a coordinate Multi-Layer Perceptron (MLP), with expressive activations [33, 28] and encodings [36] overcoming the spectral bias of early ReLU networks. To speed up the optimization process, latent feature grids [16, 22, 4, 8] reintroduced explicit, local feature storage to curb the cost of deep evaluation, retaining only a shallow decoder. Most recently, following 3D Gaussian Splatting [12], explicit primitive methods [43, 44] abandon neural decoding entirely, representing signals with continuous, freely optimizable atoms.

Contributions. In this work, we argue explicit primitives offer advantages to medical imaging research unmatched by implicit approaches, both in terms of mathematical properties and in computational abilities. Concretely: (i) we establish a unified formulation under which dictionaries, INRs, feature grids, and explicit primitives differ only in the parameterization and support of their bases; (ii) we catalogue the closed-form mathematical properties and inductive biases

that follow from explicit parametric bases, and show how each unlocks concrete advantages across established and emerging medical imaging applications; and (iii) we benchmark explicit and implicit representations across two real-world medical datasets, showing consistently higher fidelity and robustness, at lower optimization times and memory requirements.

2 Unified formulation for off-grid methods

Off-grid methods map a coordinate $\mathbf{x} \in \mathbb{R}^d$ to an intensity, feature, or vector quantity $\mathbf{v} \in \mathbb{R}^c$. We observe that the dominant paradigms are instances of a single generalized basis expansion,

$$\mathbf{v}(\mathbf{x}) = \sum_{k \in \mathcal{N}(\mathbf{x})} w_k(\mathbf{x}) \mathbf{c}_k \quad (1)$$

in which basis (weighting) functions $w_k(\mathbf{x}) \in \mathbb{R}$ are combined with coefficients $\mathbf{c}_k \in \mathbb{R}^c$ over an active set $\mathcal{N}(\mathbf{x}) \subseteq \{1, \dots, N\}$. The paradigms differ only in three choices: (i) where the bases w_k originate, (ii) their spatial support together with how the active set $\mathcal{N}(\mathbf{x})$ is selected, and (iii) which quantities are optimized.

Dictionary methods. Sparse coding approximates a signal as $\mathbf{v} \approx \sum_k \alpha_k \mathbf{d}_k$, with fixed (learned) atoms $w_k = \mathbf{d}_k$ and scalar coefficients $\mathbf{c}_k = \alpha_k$ [1]. The atoms are discrete vectors defined only at grid samples and the active set $\mathcal{N}$ is chosen *globally* by a sparsity penalty $\|\boldsymbol{\alpha}\|_0 \leq T$ rather than by spatial locality. Only the coefficients (and, in a separate stage, the dictionary) are optimized. Dictionaries are thus the explicit grid-bound ancestor of the methods below.

Implicit representations. INRs model the signal as a deep MLP $f_\theta(\mathbf{x})$ whose final linear layer reads $f_\theta(\mathbf{x}) = \sum_k h_k(\mathbf{x}) p_k$. This is exactly Eq. (1) where $w_k(\mathbf{x}) = h_k(\mathbf{x})$ is the penultimate layer’s k-th activation and $\mathbf{c} = \mathbf{p}$ are the final layer’s weights. Here, the bases are generated on-the-fly: they have global support ($\mathcal{N}(\mathbf{x}) = \{1, \dots, N\}$ for all $\mathbf{x}$), as they are nonlinear functions of *all* weights θ , and optimization targets the generator parameters θ . While early ReLU networks exhibited *spectral bias*, periodic [33] and Gabor-wavelet [28] activations and Fourier encodings [36] have shown to capture high-frequency expressivity. Generating bases on demand reduces storage requirements relative to dictionaries, although the MLP must learn to partition the domain through many successive layers, incurring large optimization times for detailed volumes.

Latent feature grids. Feature grids [16] take a hybrid approach by introducing explicit local storage alongside an implicit decoder f_θ . A (multi-resolution) grid of latent feature vectors $\mathbf{Z}$ is instantiated and queried for an individual coordinate via interpolation $\mathbf{Z}(\mathbf{x}) \in \mathbb{R}^F$. Feature grids functionally equate to explicitly representing and directly optimizing the intermediate activations $\mathbf{Z}(\mathbf{x})$ of a deep

network, bypassing the need for early layers of a network to generate it. This results in significantly lower depth requirements, considerably diminishing the compute costs of mapping coordinates to bases. Contextualizing this to fully implicit approaches, the overall function can be summarized as $\mathbf{v}(\mathbf{x}) = f_\theta(\mathbf{Z}(\mathbf{x}))$. Most famously, [22] further utilize hardware-specific tricks to improve optimization times. Alternatively, other approaches factorize dense grids into compact planes and vectors [4, 8], alleviating the storage costs of high-resolution feature grids.

Explicit representations. Explicit methods make the bases themselves the free parameters. For Gaussian primitives, we define the basis as

$$w_k(\mathbf{x}) = \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_k)^\top \boldsymbol{\Sigma}_k^{-1}(\mathbf{x} - \boldsymbol{\mu}_k)\right), \quad (2)$$

with center $\boldsymbol{\mu}_k \in \mathbb{R}^d$ and anisotropic covariance $\boldsymbol{\Sigma}_k$ parameterized through its Cholesky factors to ensure positive semi-definiteness. The Gaussian parameters $\{\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k\}$, as well as $\mathbf{c}_k$ are optimized. A Gaussian has infinite support, but its rapid exponential decay renders each primitive negligible beyond a few standard deviations. Truncating support via k-nearest neighbors further reduces computational requirements with negligible impact on reconstruction metrics.

The additive form of Eq. (1) imposes a zero (baseline) prior away from primitives. Alternatively, the normalized weighting form, known as Partition of Unity (PoU), is more efficient on large uniform regions [44], albeit forfeiting the analytic closure outlined in Section 3. This returns to the dictionary principle, yet the atoms are now continuous, may span arbitrarily-large regions, and are freely optimizable. Overlapping primitives share bases across neighboring regions, thus scaling capacity with signal complexity rather than grid resolution.

3 The Explicit Advantage: Analytic, Controllable, and Geometric Representations

Section 2 established that explicit primitives are the only paradigm whose bases are at once continuous, parametric, and geometrically explicit. We now outline existing and potential applications of these mathematical properties.

3.1 Analytic closure

Closure under convolution. Sensor blur is omnipresent in image acquisition and especially critical in medical imaging, where the measurement is the anatomy convolved with the device’s point-spread function (PSF). Modeling the PSF inside an implicit network demands costly numerical integration, propagating many stochastic samples per query [39]. Since Gaussians are closed under convolution, an explicit Gaussian representation captures the PSF as a simple covariance addition, $\boldsymbol{\Sigma}_{\text{obs}} = \boldsymbol{\Sigma}_k + \boldsymbol{\Sigma}_{\text{PSF}}$ [6]. The same rasterization extends to

tomographic projection, where it already enables sparse-view CT reconstruction an order of magnitude faster than neural fields [42].

Spatial-spectral duality. The Fourier transform of a Gaussian is again a Gaussian, so an explicit representation can be fit directly against raw frequency-domain measurements, bypassing the gridding that discrete Fourier transforms require. Gabor primitives extend Gaussian primitives by adding frequency and phase components. Their spectra are origin-shifted Gaussians, further permitting spectral space-partitioning for improved high-frequency recovery [10].

Analytical Lipschitz continuity. Balancing expressiveness and smoothness requires bounding a representation’s global upper bound on the function’s spatial rate of change, i.e. its Lipschitz constant. In implicit approaches, this demands complex processes of rigid layer-wise spectral normalizations [5], regularization techniques [15], strategic budget distribution across all network components [20], or applying smoothing constraints to the output space using control points [30]. Conversely, the Lipschitz constant of an explicit Gaussian representation is analytically defined directly by covariance matrices Σ_k and amplitudes of its primitives, entirely avoiding the architectural overhead and expressiveness bottlenecks inherent to spectrally constrained MLPs.

Closed-form derivatives and volumes. Spatial derivatives of Eq. (2) are analytic, and a primitive’s volume is proportional to $\sqrt{|\Sigma_k|}$. While strictly isochoric (volume-preserving) deformations are classically enforced through expensive Jacobian-determinant constraints [9], Gaussian representations reduce to a simple penalty on $|\Sigma_k|$, sidestepping the double-backpropagation through a neural deformation field. Explicit Gaussian deformation models already rival iterative and learning-based registration at a fraction of the runtime [14].

3.2 Controllable optimization

Structure-aware initialization. Random initialization of an MLP is notoriously hard to tune, and data-aware priors demand expensive meta-learning [35, 32]. Explicit primitives can instead be placed deterministically from the target’s own structure, concentrating modeling capacity in high-frequency regions before the first gradient step [44].

Adaptive densification. Capacity can be allocated where the signal demands it. New primitives may be spawned at high-error regions during optimization. As their support is effectively local, this resolves fine detail without perturbing already-converged regions or risking catastrophic forgetting [12]. Multi-scale primitive densification has proven effective for exactly this in deformable registration [14].

Native coarse-to-fine control. Coarse-to-fine optimization is essential to jointly learn representations and spatial transforms, as in pose or motion estimation. Implicit methods emulate it by annealing the positional encoding [25]. Meanwhile, explicit methods obtain it natively by applying a decaying blur to the covariances, granting direct control of the frequencies’ upper bound.

Inductive biases and extrapolation. The blending choice in $w(\mathbf{x})$ is itself a controllable prior. The additive form decays to a baseline away from primitives, naturally modeling empty uniform pockets (e.g. in the lungs) without the use of any primitives. Alternatively, the normalized form (PoU) captures large homogeneous regions [44, 6] with higher parameter-efficiency. Explicit primitives further extrapolate predictably beyond the training domain. This causes bordering background regions to require no additional primitives, deformation and velocity fields stay sensible outside the sampled coordinate range, and permits standardization of cohort volumes without boundary artifacts.

3.3 Geometric structure

Semantical feature extraction. Extracting semantics from an INR is not straight-forward. Mapping network weights to labels [7], meta-learning segmentation priors [38], or auto-decoding a latent vector [24] all scale poorly and tend to overfit intensities at the expense of structure [34]. Over an explicit primitive set, the representation becomes a point cloud, allowing classification and segmentation to become graph- and node-level problems addressable with permutation-invariant graph networks [26, 41]. Message passing over Gaussian primitives already drives classification [15] and segmentation [11] of splatted scenes, as well as general feature extraction [27].

Parameter-space regularization In implicit neural fields, preventing occupancy in specific regions relies on soft spatial regularizers [40, 18]. Because these are output-space constraints, they demand dense numerical integration to penalize the network indirectly. The transition to explicit parameters offers direct parameter-space intervention. If an anatomical prior dictates a disallowed region, or if a primitive’s contribution falls below a threshold, the parameters can be explicitly pruned or penalized [12, 31].

3.4 Limitations

Generative modeling. Implicit generative models rely on conditioning a multi-layer perceptron (MLP) with a global latent vector $\mathbf{z}$ [24]. Following works increase reconstruction fidelity by anchoring latent codes to specific spatial partitions [3, 21]. Some works try to bridge into the Gaussian domain by showing that Gaussian scenes can be generated by autoregressively predicting latent vectors along a grid, and learning to decode each vector into Gaussian parameters [17]. Despite the advantages of explicit methods, their off-grid topology introduces a fundamentally conflict with the strict grid-based requirements of current generative architectures. Although works exist which experiment with utilizing permutation equivariant architectures to generate Gaussian scenes [29, 13], these rely ground-truth representations trained on densely acquired data, which is a rare commodity in medical imaging. Overall, this topic remains severely under-researched, simultaneously opening an exciting frontier for modeling image distributions natively through explicit basis expansions.

4 Performance evaluation and scaling properties

Next, we evaluate the performance and scaling properties of explicit representations against common implicit representation techniques. While representation literature predominantly concerns itself in maximizing reconstruction quality, we argue that the applicability of these methods to real-world research also relies on training requirements. As such, we outline two metrics to aid researchers in quantifying hardware bottlenecks for their research:

Optimization duration. Optimization time dictates a researcher’s ability to perform architectural exploration and hyper-parameter tuning, as well as the size of a cohort for which results may be computed. We measure this as the time required for a reconstruction metric to plateau. In our experiments, we set the early-stopping criterion as the first 300 training steps elapsing without reconstruction metrics (PSNR) increasing by 0.2dB.

Backpropagation memory usage. A key factor influencing the efficacy of training processes is the size of a mini-batch. The memory capacity of a GPU places a ceiling over the maximum batch size. This influences the convergence stochasticity and the learning rate that may be employed, and consequently determines the feasibility of learning representations for a given type of data. We argue it is important for researchers to be given an up-front understanding on the memory requirements of a given method. We do so by reporting memory usage per coordinate during the backpropagation of gradients.

Datasets. To reflect applicability of implicit and explicit methods to real-world medical imaging research, we employ two datasets with significantly increased complexity compared to standard datasets used in INR literature.

Breast Cancer Histology images (BACH) [2]. Histology images offer high-resolution detailed color images of complex tissue, containing both regions of high-frequency features and flat uniform regions. The standard image size is [2048 x 1536], resulting in over 3.1 million RGB samples.

National Lung Screening Trial (NLST) [23]. 3D CT chest scans offer complex lung structures as well as intricate acquisition-related noise patterns. A volume’s resolution is [151 x 512 x 512], resulting in over 39.5 million voxels, meaning that representation methods with sensible compression factors will require to be scalable to millions of trainable parameters.

Baselines. For all methods, we tuned learning rates via grid-search to determine the training parameter choices for each method. We observed both datasets having similar optimal learning rate choices and thus use the same for both. We run all methods for compression factors 5, 10, and 20. All methods are trained using ADAM optimizer, under a batch size of 200k and mixed-precision floating-point assignments.

Gaussian methods. We evaluate two Gaussian weighting schemes: Additive [43] and PoU [44]. We follow the initialization scheme in [44] and use a learning rate of $5 \cdot 10^{-2}$.




Fig. 2. Performance and efficiency metrics for various implicit and explicit approaches. Scatter point grayness indicates the number of steps before a run failed to increase PSNR by 0.2 and was early-stopped (black=earlier, white=later).

Implicit representations. We choose three popular fully-implicit baselines: Fourier-features [36], SIREN [33], and WIRE [28]. We used a learning rate of 10^{-4} for Fourier-features, while SIREN and WIRE appeared to require a learning rate of 10^{-5} in order to achieve competitive results. For Fourier-features we used 256 frequencies under the BACH dataset and 512 for NLST. For SIREN a frequency term ω of 250 performed best for BACH while a lower ω of 30 performed best for NLST. Similarly, WIRE performed best with parameters $\omega = 250$ and $\sigma = 0.1$ for BACH, while under NLST WIRE performed best with $\omega = 10$ and $\sigma = 1$.

Feature grids. We include two feature-grid approaches: Multi-resolution Neural Volumes [16] and K-planes [8]. Both feature grids resolutions start from 16 and increase by a factor of 2 up until the compression quota is met (resulting in the final resolution being a partial factor of 2). A learning rate of 10^{-3} was used for both Neural Volumes and K-planes.

Results. Gaussian primitives appear to match or surpass implicit baselines in PSNR metrics. The Gaussian PoU weighting scheme seems to surpass the plain additive scheme, likely to its weighting scheme’s improved ability to model uniform regions more efficiently. While compared to feature grids, Gaussians are 2-5x faster and require 5-15x less memory, the difference grows to 5-100x faster and 10-200x less memory when compared to fully implicit methods. Fully-implicit methods scale the poorest, requiring the most fitting time and memory, while achieving lower metrics than Gaussian PoU and Neural Volumes. Also, as seen in the 8-layer variants in Fig. 2, fully-implicit methods suffer under increased depths likely due to vanishing gradients.




Fig. 3. Performance and efficiency metrics for various implicit and explicit approaches. Scatter point grayness indicates the number of steps before a run failed to increase PSNR by 0.2 and was early-stopped (black=earlier, white=later).

5 Conclusion

In this work we have consolidated implicit and explicit representations into the same mathematical framework of dictionary methods. Furthermore, we have outlined the various mathematical properties that, from a medical imaging perspective, make Gaussian primitives exceptionally interesting over implicit approaches. Finally, our empirical results on two datasets show that explicit methods offer the most scalable solution to image representation while matching or surpassing state-of-the-art implicit methods. Ultimately, we hope that this work will serve as a catalyst to inspire a new generation of efficient and scalable off-grid paradigm.

Acknowledgments

We would like to thank Julian McGinnis for his endless patience with our continuous Gaussian discussions despite his implicit preferences. This work is funded by the Munich Center for Machine Learning.

Disclosure of Interests

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

References

1. Aharon, M., Elad, M., Bruckstein, A.: K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Trans. Signal Process.* **54**(11), 4311–4322 (2006)
2. Aresta, G., et al.: BACH: Grand challenge on breast cancer histology images. *Med. Image Anal.* **56**, 122–139 (2019)
3. Chabra, R., Lenssen, J.E., Ilg, E., Schmidt, T., Straub, J., Lovegrove, S., Newcombe, R.: Deep local shapes: Learning local SDF priors for detailed 3d reconstruction. In: *ECCV* (2020)
4. Chen, A., Xu, Z., Geiger, A., Yu, J., Su, H.: TensorRF: Tensorial radiance fields. In: *ECCV* (2022)
5. Coiffier, G., Béthune, L.: 1-Lipschitz neural distance fields. *Comput. Graph. Forum* **43**(5), e15128 (2024)
6. Dannecker, M., et al.: Fast and explicit: Slice-to-volume reconstruction via 3d Gaussian primitives with analytic point spread function modeling. *arXiv preprint arXiv:2512.11624* (2025)
7. De Luigi, L., Cardace, A., Spezialetti, R., Zama Ramirez, P., Salti, S., Di Stefano, L.: Deep learning on implicit neural representations of shapes. In: *ICLR* (2023)
8. Fridovich-Keil, S., Meanti, G., Warburg, F.R., Recht, B., Kanazawa, A.: K-Planes: Explicit radiance fields in space, time, and appearance. In: *CVPR* (2023)
9. Haber, E., Modersitzki, J.: Numerical methods for volume preserving image registration. *Inverse Probl.* **20**(5), 1621–1638 (2004)
10. Huang, W., et al.: Gabor primitives for accelerated cardiac cine MRI reconstruction. *arXiv preprint arXiv:2603.05681* (2026)
11. Jain, U., Mirzaei, A., Gilitschenski, I.: GaussianCut: Interactive segmentation via graph cut for 3d Gaussian splatting. In: *NeurIPS* (2024)
12. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d Gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.* **42**(4), 1–14 (2023)
13. Lan, Y., Zhou, S., Lyu, Z., Hong, F., Yang, S., Dai, B., Pan, X., Loy, C.C.: GaussianAnything: Interactive point cloud flow matching for 3d generation. In: *ICLR* (2025)
14. Li, X., et al.: Gaussian primitives for deformable image registration. *arXiv preprint arXiv:2406.03394* (2024)
15. Liu, H.T.D., Williams, F., Jacobson, A., Fidler, S., Litany, O.: Learning smooth neural functions via Lipschitz regularization. In: *ACM SIGGRAPH* (2022)
16. Lombardi, S., Simon, T., Saragih, J., Schwartz, G., Lehmman, A., Sheikh, Y.: Neural volumes: Learning dynamic renderable volumes from images. *arXiv preprint arXiv:1906.07751* (2019)
17. von Lütow, N., Röfle, B., Schmid, K., Nießner, M.: GaussianGPT: Towards autoregressive 3d Gaussian scene generation. *arXiv preprint arXiv:2603.26661* (2026)
18. Maas, K.W., Ruijters, D., Vilanova, A., Pezzotti, N.: NeRF-CA: Dynamic reconstruction of x-ray coronary angiography with extremely sparse views. *IEEE Trans. Vis. Comput. Graph.* (2025)
19. Mairal, J., Bach, F., Ponce, J., Sapiro, G.: Online dictionary learning for sparse coding. In: *ICML* (2009)
20. McGinnis, J., Shit, S., Hölzl, F.A., Friedrich, P., Büschl, P., Sideri-Lampretsa, V., Mühlau, M., Cattin, P.C., Menze, B., Rueckert, D., et al.: Beyond uniformity: Regularizing implicit neural representations through a Lipschitz lens. In: *ICLR* (2026)

21. Mehta, I., Gharbi, M., Barnes, C., Shechtman, E., Ramamoorthi, R., Chandraker, M.: Modulated periodic activations for generalizable local functional representations. In: ICCV (2021)
22. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.* **41**(4), 1–15 (2022)
23. National Lung Screening Trial Research Team: The national lung screening trial: Overview and study design. *Radiology* **258**(1), 243–253 (2011)
24. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: DeepSDF: Learning continuous signed distance functions for shape representation. In: CVPR (2019)
25. Park, K., Sinha, U., Barron, J.T., Bouaziz, S., Goldman, D.B., Seitz, S.M., Martin-Brualla, R.: Nerfies: Deformable neural radiance fields. In: ICCV (2021)
26. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: PointNet: Deep learning on point sets for 3d classification and segmentation. In: CVPR (2017)
27. Qi, Z., Ma, J., Xu, J., Zhou, Z., Cheng, L., Xiong, G.: GSPR: Multimodal place recognition using 3d Gaussian splatting for autonomous driving. In: IROS (2025)
28. Saragadam, V., LeJeune, D., Tan, J., Balakrishnan, G., Veeraraghavan, A., Baraniuk, R.G.: WIRE: Wavelet implicit neural representations. In: CVPR (2023)
29. Shabanov, A., Hedman, P., Weber, E., Li, Z., Rozumny, D., Lan, G.L., Dhingra, N., Luo, L., Vedaldi, A., Richardt, C., et al.: Free-range gaussians: Non-grid-aligned generative 3d gaussian reconstruction. *arXiv preprint arXiv:2604.04874* (2026)
30. Sideri-Lampretsa, V., McGinnis, J., Qiu, H., Paschali, M., Simson, W., Rueckert, D.: SINR: Spline-enhanced implicit neural representation for multi-modal registration. In: MIDL (2024)
31. 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 (2026)
32. Sitzmann, V., Chan, E.R., Tucker, R., Snavely, N., Wetzstein, G.: MetaSDF: Meta-learning signed distance functions. In: NeurIPS (2020)
33. Sitzmann, V., Martel, J.N., Bergman, A.W., Lindell, D.B., Wetzstein, G.: Implicit neural representations with periodic activation functions. In: NeurIPS (2020)
34. Stolt-Ansó, N., McGinnis, J., Pan, J., Hammernik, K., Rueckert, D.: NISF: Neural implicit segmentation functions. In: MICCAI (2023)
35. Tancik, M., Mildenhall, B., Wang, T., Schmidt, D., Srinivasan, P.P., Barron, J.T., Ng, R.: Learned initializations for optimizing coordinate-based neural representations. In: CVPR (2021)
36. Tancik, M., Srinivasan, P.P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J.T., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains. In: NeurIPS (2020)
37. Van Nguyen, H., Patel, V.M., Nasrabadi, N.M., Chellappa, R.: Kernel dictionary learning. In: ICASSP (2012)
38. Vyas, K., Veeraraghavan, A., Balakrishnan, G.: Fit pixels, get labels: Meta-learned implicit networks for image segmentation. In: MICCAI (2025)
39. Xu, J., Moyer, D., Gagoski, B., Iglesias, J.E., Grant, P.E., Golland, P., Adalsteinsson, E.: NeSVoR: Implicit neural representation for slice-to-volume reconstruction in MRI. *IEEE Trans. Med. Imaging* **42**(6), 1707–1719 (2023)
40. Yang, J., Pavone, M., Wang, Y.: FreeNeRF: Improving few-shot neural rendering with free frequency regularization. In: CVPR (2023)
41. Zaheer, M., Kottur, S., Ravanbakhsh, S., Póczos, B., Salakhutdinov, R., Smola, A.J.: Deep sets. In: NeurIPS (2017)
42. Zha, R., Lin, T.J., Cai, Y., Cao, J., Zhang, Y., Li, H.: R²-Gaussian: Rectifying radiative Gaussian splatting for tomographic reconstruction. In: NeurIPS (2024)

- 43. Zhang, X., Ge, X., Xu, T., He, D., Wang, Y., Qin, H., Lu, G., Geng, J., Zhang, J.: GaussianImage: 1000 FPS image representation and compression by 2d Gaussian splatting. In: ECCV (2024)
- 44. Zhang, Y., Li, B., Kuznetsov, A., Jindal, A., Diolatzis, S., Chen, K., Sochenov, A., Kaplanyan, A., Sun, Q.: Image-GS: Content-adaptive image representation via 2d Gaussians. In: ACM SIGGRAPH (2025)