2026 JULY 19–22 VANCOUVER, BC
JOINT AAPM | COMP MEETING

DiffINR: Conditional Neural Field Diffusion for Deformation-Driven Volumetric Image Estimation from Limited Measurements

Xiaoxue Qian[1,2], Hua-Chieh Shao[1,2], You Zhang[1,2]
[1]Department of Radiation Oncology, UT Southwestern Medical Center, Dallas, Texas
[2]Medical Artificial Intelligence and Automation (MAIA) Lab

MAIA Laboratory
UT Southwestern Medical Center
Radiation Oncology

ABSTRACT

Rapid and accurate estimation of patient-specific volumetric images from limited measurements remains challenging in time-sensitive medical imaging applications. Coordinate-based implicit neural representations (INRs) have shown strong potential for high-quality volumetric modeling; however, optimizing an INR from scratch for each patient is computationally costly and prone to suboptimal convergence under severe information scarcity. In this study, we propose DiffINR, a conditional hyperdiffusion framework that operates directly in INR weight space to learn the distribution of implicit neural fields, enabling efficient volumetric estimation. It predicts patient-specific INR initialization for deformation-driven volumetric image estimation from limited measurements. First, a collection of INRs was optimized to represent individual samples from fully sampled training data, and the resulting weights were used as ground-truth signals for diffusion model training. Next, a conditional hyperdiffusion model was trained in the INR parameter space to predict INR weights using two complementary conditioning features: foundation model-derived semantic features extracted from paired source/target images, and measurement-aware features encoded from available measurements to preserve patient specificity. Finally, the predicted weights served as informative initialization for rapid test-time refinement, where the INR predicts DVFs to deform a high-quality prior source image for volumetric image estimation. Experiments were conducted on undersampled volumetric MR imaging and limited-angle CBCT estimation tasks. DiffINR achieved competitive or superior performance across imaging tasks, with improved image quality, registration accuracy, and robustness under increasingly severe data limitations, while reducing test-time iterations and runtime. As a neural-field generative framework, DiffINR enables efficient, high-quality 3D image estimation from undersampled MR k-space and limited-angle CBCT projections by combining population priors with rapid patient-specific adaptation, supporting real-time image guidance, motion management, and adaptive radiotherapy workflows.

CONTACT

Dr. You Zhang:
You.Zhang@UTSouthwestern.edu

INTRODUCTION

Volumetric imaging plays a critical role in image-guided radiotherapy by providing patient-specific anatomical information for target localization, motion monitoring, and adaptive treatment. However, accelerating image acquisition often results in limited-angle CBCT projections or undersampled MR k-space, making accurate volumetric estimation highly challenging. Although learning-based deformation methods enable fast inference, their performance relies on population-level mappings and may generalize poorly to unseen patients. Coordinate-based implicit neural representations (INRs) provide flexible patient-specific optimization but require expensive and unstable optimization from scratch. To address these limitations, we propose DiffINR, a conditional hyperdiffusion framework that predicts informative INR initialization in parameter space for deformation-driven volumetric image estimation, where a high-quality prior source image is deformed to new target images using limited measurements.

METHODS

The proposed DiffINR consists of three stages (Fig. 1).

(1) INR Parameter Space Construction.
For a training cohort of fully sampled volumetric image pairs, patient-specific INRs are optimized to represent deformation vector fields (DVFs), resulting in optimized INR parameters that characterize population-level deformation patterns. These parameters are used to construct the INR parameter space for subsequent diffusion learning.

(2) Conditional Hyperdiffusion Learning.
Based on Hyperdiffusion [1], a conditional diffusion model is trained in the INR parameter space to learn the distribution of optimized INR parameters using foundation model-derived semantic priors and measurement-aware features. The trained model predicts informative INR initialization for patient-specific adaptation.

(3) Test-time Patient-specific Refinement.
For a new test case, the trained diffusion model predicts patient-specific INR parameters from available measurements to initialize rapid test-time refinement. The initialized INR is then optimized under the measurement fidelity objective for deformation-driven volumetric image estimation.

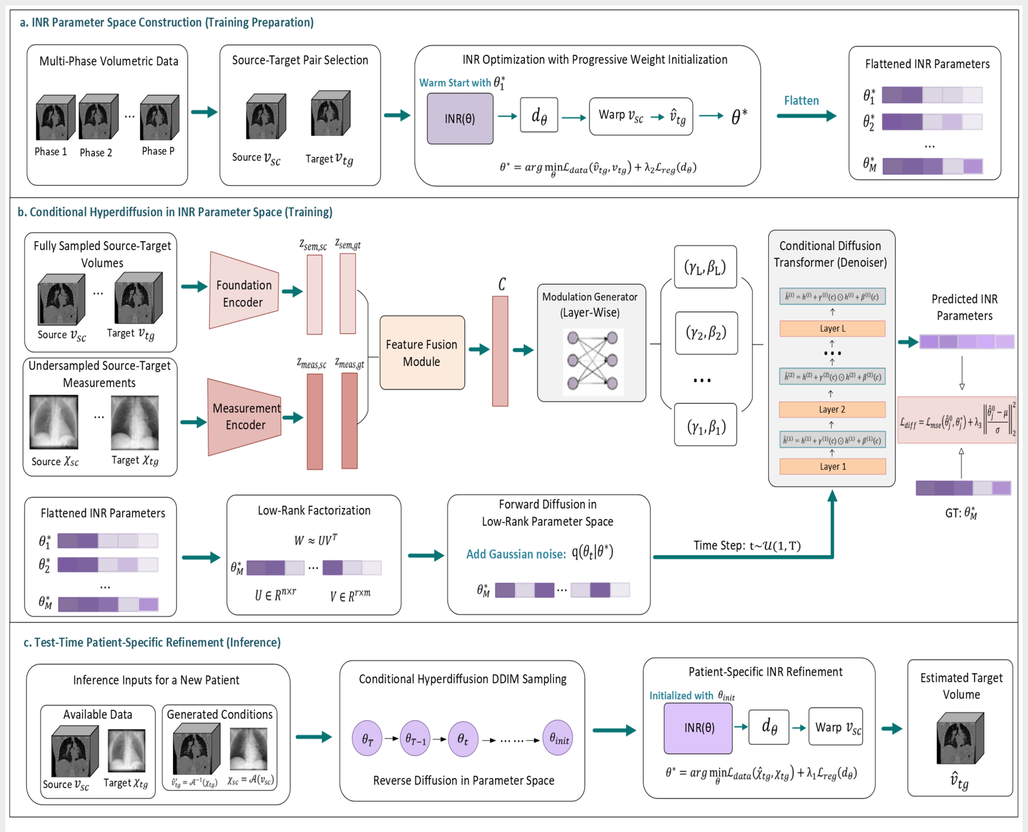

Figure 1. Overview of the proposed DiffINR framework.

RESULTS

DiffINR was evaluated for limited-angle CBCT estimation, where 26 lung 4D-CT cases (TCIA and CREATIS) were used for training and 10 DIR-Lab cases for testing. Limited-angle CBCT was estimated by deforming prior CT images. DiffINR was compared with eight methods, including conventional iterative optimization, population-based registration networks, patient-specific INR optimization, and hybrid population-informed INR approaches [2,3]. Qualitative results and quantitative comparisons are shown in Fig. 2 and Table 1, respectively. DiffINR was also evaluated on the OCMR cardiac MRI dataset for real-time volumetric MRI estimation, using 5-fold cross-validation, with 9 subjects for training and 2 subjects for testing in each fold. Real-time volumetric MRIs estimated from deforming prior MRIs, as guided by undersampled k-space measurements, were compared between nine methods. Results are presented in Fig. 3 and Table 2.

Figure 2. Comparison of estimated CBCTs and corresponding difference maps for one case from the DIR-Lab dataset.

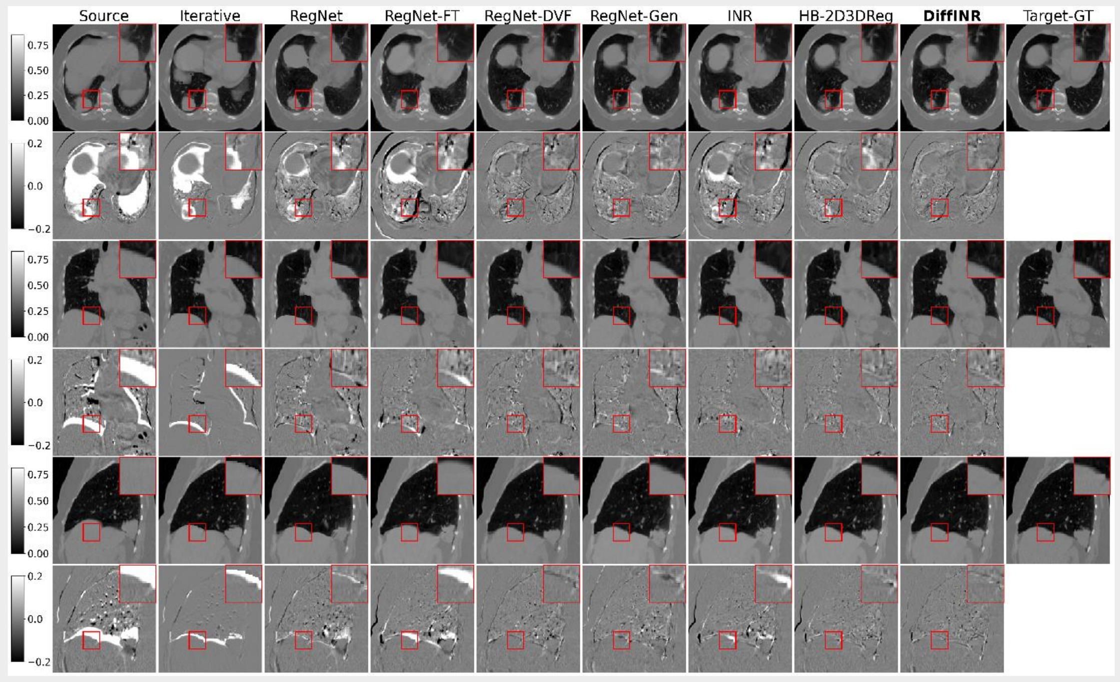

Table 1. Comparison of SSIM, RE, and TRE metrics across nine methods on the DIR-Lab testing dataset.

Methods / DIR-Lab	Mean (± s.d.)		
	SSIM↑	RE/%↓	TRE/mm↓
No Registration	0.925 ± 0.035	15.40 ± 2.41	8.52 ± 3.31
Iterative	0.946 ± 0.040	8.85 ± 2.34	5.02 ± 2.37
RegNet	0.942 ± 0.034	9.82 ± 2.12	6.38 ± 2.46
RegNet-FT	0.950 ± 0.024	8.73 ± 1.98	4.94 ± 2.35
RegNet-DVF	0.945 ± 0.035	9.02 ± 2.62	5.35 ± 2.72
RegNet-Gen	0.948 ± 0.031	8.97 ± 2.54	5.11 ± 2.13
INR	0.949 ± 0.035	9.71 ± 2.33	5.01 ± 2.77
HB-2D3DReg	0.953 ± 0.035	7.99 ± 2.16	3.70 ± 1.94
DiffINR	**0.962 ± 0.021**	**6.84 ± 1.70**	**3.61 ± 1.27**

Figure 3. Comparison of estimated MRIs and corresponding difference maps for one case from the OCMR dataset.

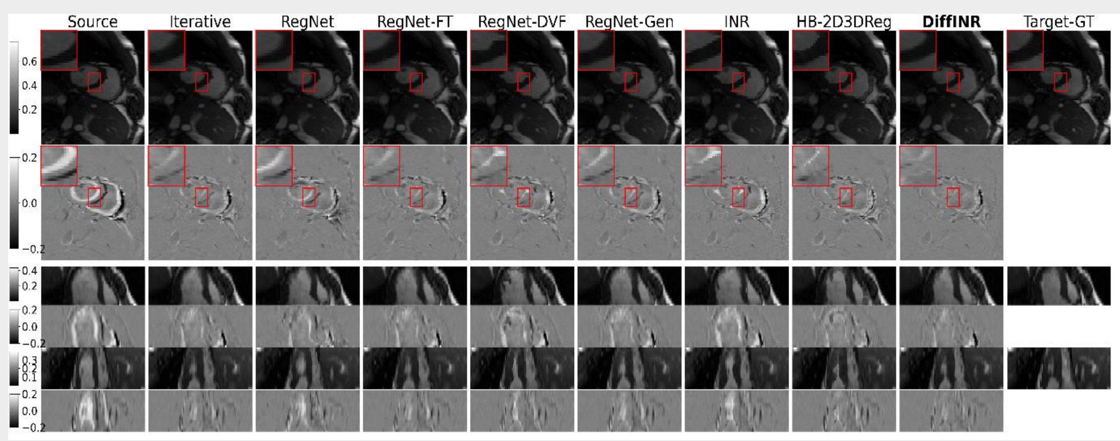

Table 2. Comparison of SSIM, RE, Dice, HD95, and COME across nine methods on the OCMR dataset.

Methods /OCMR	Mean (± s.d.)				
	SSIM↑	RE/%↓	Dice↑	HD95↓	COME/mm↓
No Registration	0.956 ± 0.007	13.45 ± 4.46	0.77 ± 0.12	6.24 ± 2.58	2.47 ± 1.66
Iterative	0.979 ± 0.009	8.86 ± 4.13	0.86 ± 0.02	4.54 ± 1.82	1.57 ± 1.28
RegNet	0.976 ± 0.007	9.64 ± 4.21	0.87 ± 0.05	4.65 ± 1.48	1.71 ± 1.33
RegNet-FT	0.979 ± 0.009	9.00 ± 3.57	0.87 ± 0.07	4.53 ± 1.92	1.58 ± 1.25
RegNet-DVF	0.972 ± 0.006	8.34 ± 4.37	0.85 ± 0.06	5.01 ± 2.05	1.83 ± 1.53
RegNet-Gen	0.978 ± 0.007	8.26 ± 4.12	0.87 ± 0.06	4.49 ± 1.38	1.50 ± 1.06
INR	0.976 ± 0.008	8.78 ± 3.28	0.86 ± 0.05	4.72 ± 1.53	1.74 ± 1.18
cMeta-INR	0.983 ± 0.007	**7.44 ± 2.76**	0.89 ± 0.07	**3.53 ± 1.26**	1.32 ± 1.02
DiffINR	**0.985 ± 0.007**	8.03 ± 2.16	**0.89 ± 0.04**	4.06 ± 1.33	**1.30 ± 1.09**

DISCUSSION

The proposed DiffINR demonstrated consistent performance on both limited-angle CBCT and undersampled volumetric MRI estimation, indicating that the learned INR parameter distribution provides effective initialization for patient-specific optimization across different imaging tasks. Compared with conventional iterative optimization, population-based learning methods, and existing INR-based approaches, DiffINR combines population-level deformation priors with test-time patient-specific adaptation, enabling efficient optimization while maintaining the flexibility of subject-specific modeling. The qualitative and quantitative results on both the DIR-Lab and OCMR datasets further demonstrate that diffusion-based parameter initialization improves anatomical consistency and deformation accuracy under limited measurements. These findings suggest that conditional diffusion in INR parameter space provides an effective strategy for deformation-driven volumetric imaging across different modalities and acquisition settings.
Future work will evaluate DiffINR on a broader range of limited-data volumetric imaging tasks to further assess its robustness and generalizability across different imaging modalities, anatomical sites, and acquisition settings. We will also investigate strategies to further improve the efficiency of diffusion-based parameter initialization and patient-specific refinement, with the goal of achieving sub-second test-time computation while maintaining estimation accuracy. Continued improvements in both generalizability and computational efficiency may facilitate broader clinical applications in real-time image guidance, motion management, adaptive radiotherapy, and other image-guided radiotherapy applications.

CONCLUSION

In this study, we developed DiffINR, a conditional neural-field diffusion framework that generates patient-specific INR initialization parameters in the neural field weight space, for deformation-driven volumetric image estimation from limited measurements. By combining conditional diffusion learning with subsequent patient-specific refinement, DiffINR achieved accurate and efficient volumetric image estimation. Experimental results demonstrated that DiffINR improved accuracy, robustness, and optimization efficiency, supporting its potential for rapid volumetric imaging in image-guided and adaptive radiotherapy workflows.

REFERENCES

[1] Erkoç, Z., et al. *Hyperdiffusion: Generating implicit neural fields with weight-space diffusion.* in *Proceedings of the IEEE/CVF international conference on computer vision*. 2023.
[2] Qian, X., H.C. Shao, and Y. Zhang, *A hybrid population-based and patient-specific framework for 2D-3D deformable registration-driven limited-angle cone-beam CT estimation.* Medical physics, 2025. **52**(11): p. e70098
[3] Qian, X., et al., *cMeta-INR: cohort-informed meta-learning-based implicit neural representation for deformable registration-driven real-time volumetric MRI estimation.* Physics in Medicine & Biology, 2026. **71**(1): p. 015003.