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

One-Shot Physics-Informed Implicit Neural Representation for Low-Dose Dynamic Myocardial Perfusion CT Reconstruction

Yu Gao, Tingliang Zhuang, You Zhang
Department of Radiation Oncology, UT Southwestern Medical Center, Dallas, TX

ABSTRACT

Purpose. Dynamic myocardial perfusion CT (MPCT) is a critical tool for functional assessment of coronary stenosis and myocardial ischemia, but the continuous acquisition needed to track contrast kinetics delivers a substantial radiation dose. Low-mAs protocols reduce dose yet suffer severe quantum noise and streak artifacts that degrade spatiotemporal fidelity and prevent accurate quantitative perfusion maps. We propose MPCT-INR to recover anatomical structure and dynamic contrast from low-dose projection data.

Methods. An Implicit Neural Representation (INR) tightly models the spatiotemporal dynamic CT sequence. A physics-informed, two-stage one-shot optimization first pre-trains on noisy per-frame FBP reconstructions for a stable initialization, then jointly optimizes a composite loss of projection-domain data fidelity plus spatial-smoothness and temporal-continuity regularization.

Results. On low-dose XCAT simulations MPCT-INR delivered a 20.1% PSNR gain and 58.1% lower relative error versus FBP, and stayed stable where an explicit tensor baseline showed semi-convergence and noise overfitting. Time-attenuation curves for ventricles and myocardium captured wash-in / wash-out dynamics with high temporal fidelity.

Conclusion. MPCT-INR reconstructs anatomy and recovers dynamic contrast from low-dose data, outperforming baselines and showing great promise for dose-efficient dynamic myocardial perfusion imaging.

KEY CONTRIBUTIONS

- Continuous spatiotemporal INR (hash-encoded MLP) replaces frame-wise voxel grids, sharing information across all time frames.
- Physics-informed, self-supervised one-shot training, no paired/clean training data required.
- Stable optimization with a broad early-stopping window, avoiding the semi-convergence of explicit methods.
- +20.1% PSNR and -58.1% relative error vs FBP on low-dose XCAT MPCT.

CONTACT

Yu Gao, PhD
UT Southwestern Medical Center, Dallas, TX
Yu.Gao@UTSouthwestern.edu
Poster Board: SU-315-GPD-D-111

INTRODUCTION

- Dynamic MPCT is a critical tool for the functional assessment of coronary artery disease and myocardial ischemia, but the continuous acquisition required to track contrast kinetics delivers a substantial radiation dose to the patient.
- Low-mAs scans reduce exposure but introduce severe quantum noise and streak artifacts that degrade spatiotemporal fidelity and compromise the quantification of hemodynamic parameters.
- Existing reconstruction methods struggle to balance noise suppression against the preservation of rapid temporal dynamics.
- We propose MPCT-INR: unlike discrete voxel-based methods, the Implicit Neural Representation parameterizes the spatiotemporal images as a continuous function, naturally enforcing spatial and temporal coherence and recovering high-fidelity dynamic sequences from noisy projections.

METHODS

- **Continuous representation:** a coordinate-based MLP equipped with a multiresolution Hash Encoder maps each spatiotemporal coordinate $(\boldsymbol{x}, t)$ to its linear attenuation coefficient, so the whole dynamic sequence is one continuous function $f(\boldsymbol{x}, t)$ rather than a stack of independent voxel frames, sharing information across all time points.
- **Two-stage one-shot training:** Stage 1 pre-trains the network on noise-corrupted per-frame FBP reconstructions to give a stable initialization; Stage 2 jointly optimizes all frames against the measured low-dose projections, self-supervised, with no clean or paired training data.
- **Composite loss:** projection-domain data fidelity (forward-projected INR vs. measured data) + spatial total variation (TV) to suppress quantum noise + temporal TV to preserve the wash-in / wash-out continuity

$$\mathcal{L} = \sum_i ||\mathcal{P}_i(f(\boldsymbol{x}, t_i)) - \boldsymbol{y}_i||_1 + \lambda_{spa} \sum_i ||\nabla_x f(\boldsymbol{x}, t_i)||_1 + \lambda_{temp} \sum_i ||\partial_t f(\boldsymbol{x}, t_i)||_1.$$

- **Baselines (identical loss & FBP init):** (1) FBP per frame, standard analytical reconstruction; (2) Explicit tensor, the image as a discrete voxel grid optimized directly through the PyTorch backend, isolating the benefit of the neural representation.
- **Phantom & dynamics:** 2D low-dose fan-beam XCAT phantom; ground-truth contrast dynamics synthesized with the standard MPCT convolution model, tissue residue functions parameterized by Fermi functions to mimic realistic wash-in / wash-out kinetics.
- **Acquisition & noise:** per frame, 180 projection views over a full rotation with a 512-bin detector (0.2 cm bins); images reconstructed on a 256x256 grid at 0.16 cm isotropic voxels; low incident photon intensity with Poisson noise plus Gaussian electronic noise (variance 10).
- **Evaluation:** visual quality, time-attenuation curves (TACs), and PSNR / SSIM / Relative Error, assessed after 2000 training epochs with convergence tracked out to 10,000 epochs.

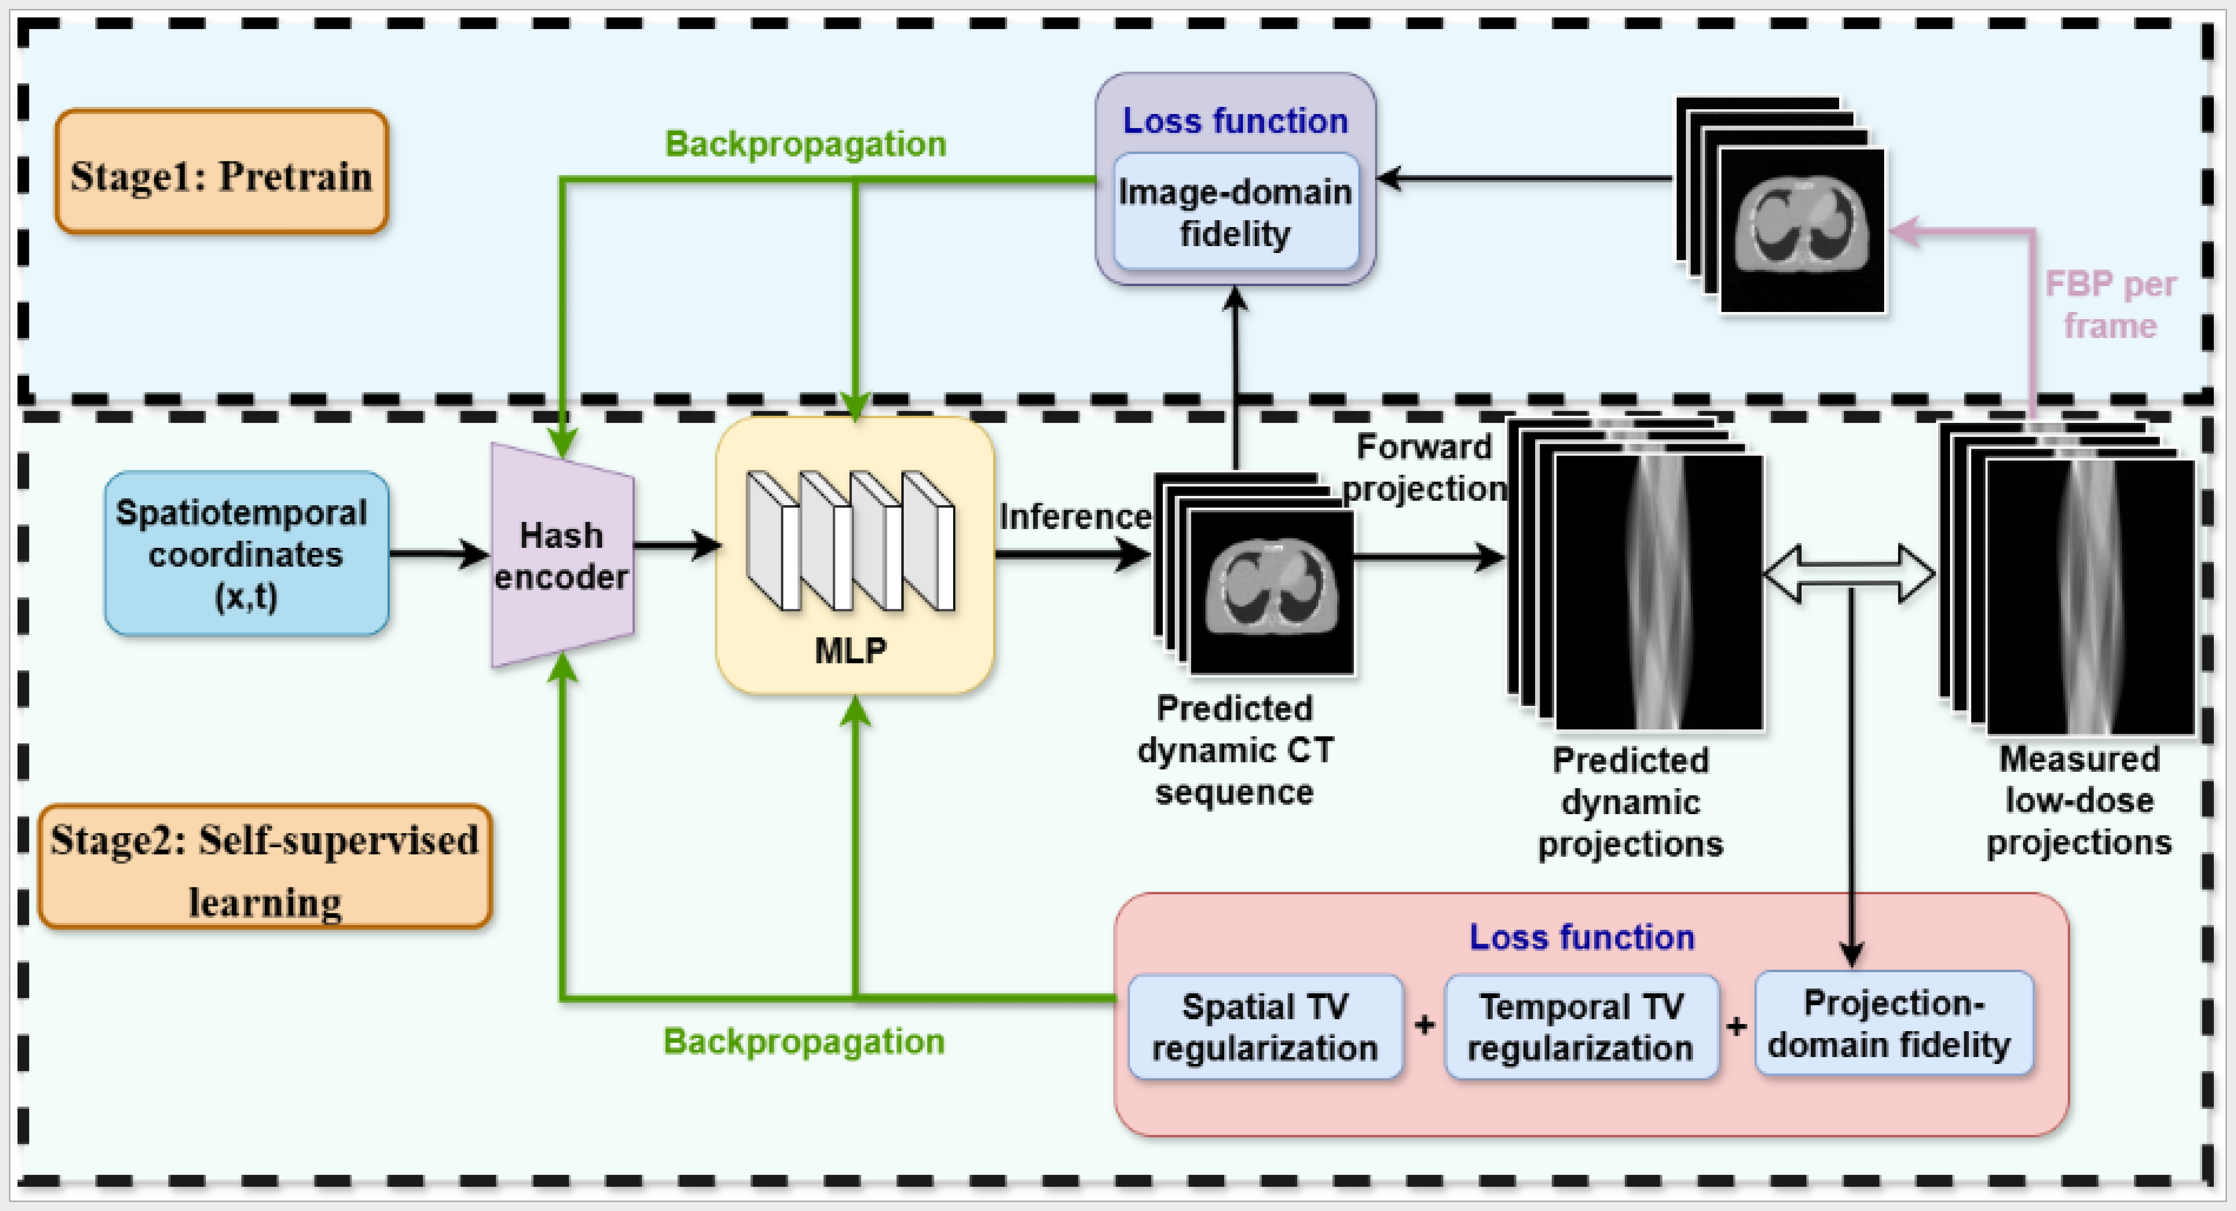

Fig. 1. *MPCT-INR pipeline: Stage 1 pre-trains on per-frame FBP reconstructions; Stage 2 jointly optimizes all frames against the measured low-dose projections (self-supervised, projection-domain).*

RESULTS

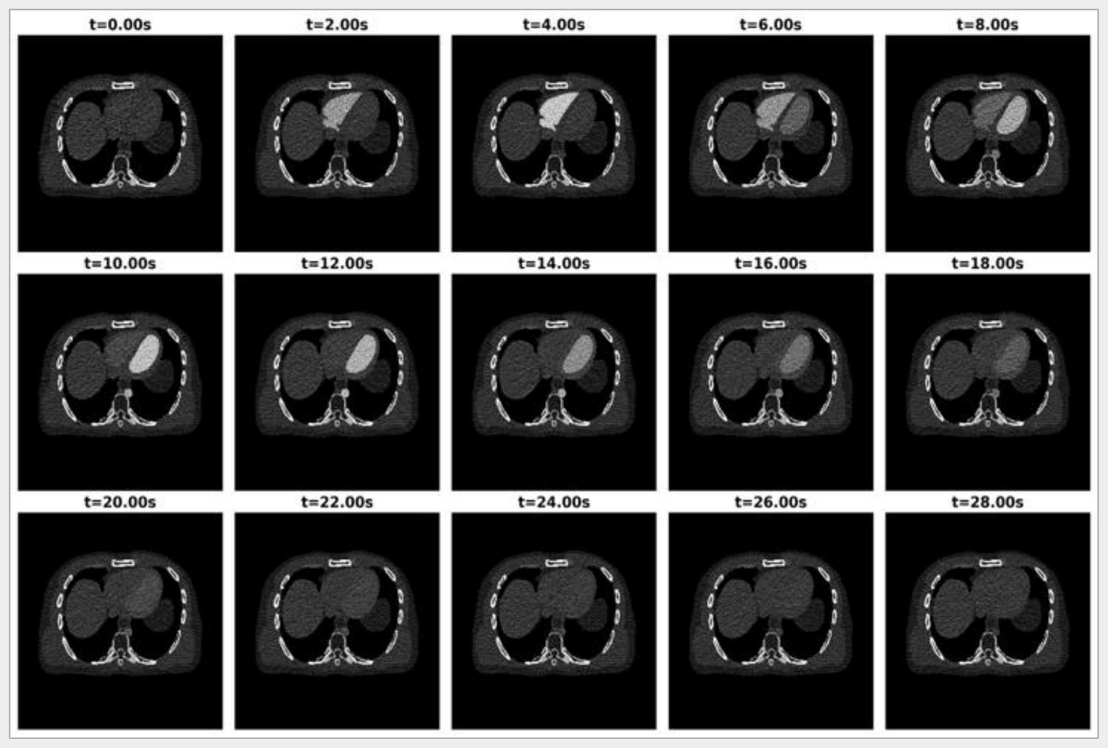

Fig. 2. *Temporal contrast dynamics (wash-in / wash-out) resolved by MPCT-INR from low-dose data across the cardiac cycle.*

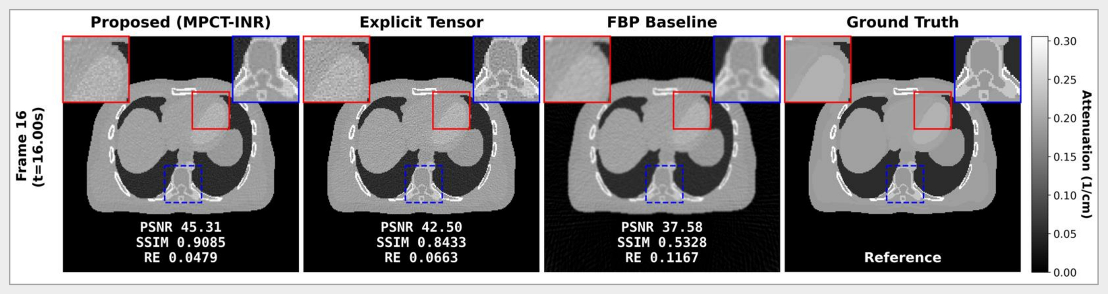

Fig. 3. *Reconstruction at t = 16 s, MPCT-INR, explicit tensor, FBP, ground truth. MPCT-INR best recovers boundaries & myocardial texture.*

- MPCT-INR achieves the best contrast and structural clarity; FBP is blurred with insufficient contrast, while the explicit tensor is corrupted by severe noise.

Table 1. *Quantitative metrics at t = 16 s.*

Method	PSNR	SSIM	Rel. Err
MPCT-INR	**45.31**	**0.9855**	**0.0479**
Explicit tensor	42.50	0.8433	0.0663
FBP	37.58	0.5328	0.1167

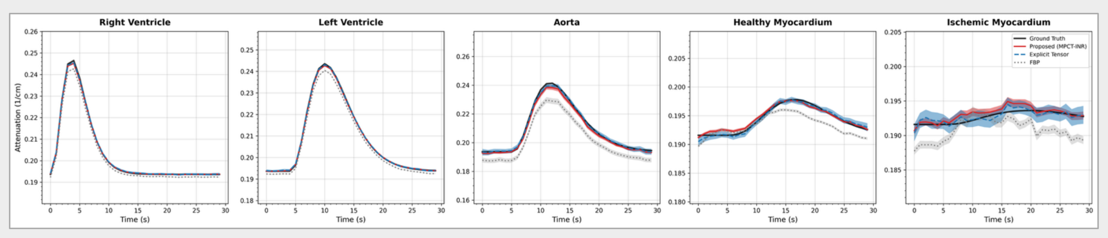

Fig. 4. *Time-attenuation curves (RV, LV, aorta, healthy & ischemic myocardium); MPCT-INR tracks ground truth with narrow SD bands.*

DISCUSSION

- Across all tissues MPCT-INR tracks the ground truth, capturing rapid wash-in peaks and wash-out tails; narrow SD bands indicate high spatial stability and signal homogeneity.
- FBP is unstable and biased in the healthy and ischemic myocardium, making quantification unreliable; the explicit tensor follows the mean trend but has much wider error bands, especially in ischemic myocardium.

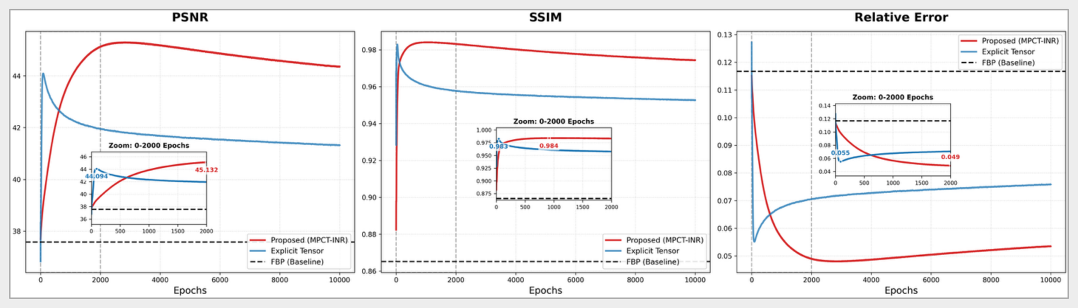

Fig. 5. *Convergence of PSNR / SSIM / Relative Error over 10,000 epochs (insets: epochs 0-2000).*

- The explicit tensor shows clear semi-convergence, a narrow optimality window that complicates early stopping. MPCT-INR degrades far more slowly, enabling a robust, practical early-stopping strategy.
- Robustness stems from the INR's continuous architecture and shared Hash encoding, which act as strong spatiotemporal regularizers isolating true dynamic contrast from projection noise.

CONCLUSIONS

- MPCT-INR effectively reconstructs anatomy and recovers dynamic contrast from noisy low-dose projections, outperforming FBP and explicit tensor optimization.
- Implicit continuous representation + shared hash encoding give strong spatiotemporal regularization and stable optimization.
- Great promise for dose-efficient dynamic myocardial perfusion CT in clinical applications.

REFERENCES

1. Segars WP, Sturgeon G, Mendonca S, Grimes J, Tsui BM. 4D XCAT phantom for multimodality imaging research. Med Phys. 2010;37(9):4902–4915.
2. Shen L, Pauly J, Xing L. NeRP: implicit neural representation learning with prior embedding for sparsely sampled image reconstruction. IEEE Trans Neural Netw Learn Syst. 2022;35(1):770–782.
3. Müller T, Evans A, Schied C, Keller A. Instant neural graphics primitives with a multiresolution hash encoding. ACM Trans Graph. 2022;41(4):102.