Hey all,

Here is what's been going on in medical physics and related areas the last 24 hours. Here's what you folks might find interesting today.

Imaging Physics and Reconstruction / Novel Architectures for Imaging

Harnessing Magnitude-Only and Complex Measurements for Improved Dynamic MRI Reconstruction with Learned Priors introduces a genuinely fresh angle on MRI reconstruction by incorporating auxiliary k-space magnitude information alongside standard complex-valued measurements. The key insight is that k-space magnitudes exhibit strong temporal consistency in steady-state dynamic MRI, providing complementary constraints that conventional physics-driven deep learning methods ignore. The authors embed this into an ADMM-unrolled framework with a novel magnitude-aware data-fidelity term, tackling the resulting non-convexity and non-differentiability via quadratically smoothed optimization and momentum-based updates. Validation spans retrospectively undersampled cine and phase-contrast flow MRI as well as prospective real-time acquisitions, with blinded reader studies confirming improved artifact suppression and phase preservation. This matters because it expands the information-theoretic basis of reconstruction beyond the standard complex k-space paradigm without adding scan time.

Primitive Representation Learning for Unsupervised Dynamic Contrast Enhanced MRI Reconstruction extends scan-specific Gaussian and Gabor primitive representations -- previously used for static MRI -- into the dynamic contrast-enhanced (DCE) setting. The architecture disentangles underlying anatomy, dynamic contrast enhancement, and residual motion into separate temporal basis functions, giving a geometrical interpretation to the learned representation. This is notable because it avoids large training datasets entirely (unsupervised, scan-specific) while achieving competitive reconstruction quality and accurate pharmacokinetic curve extraction for aorta and kidney DCE-MRI. The modular tier design is extensible to additional dynamic factors and higher acceleration rates, making it a principled alternative to data-hungry supervised approaches for quantitative DCE analysis.

Learning Where and What to Lift for Bi-planar X-ray-to-CT Reconstruction tackles the severely ill-posed bi-planar X-ray-to-CT problem by explicitly decoupling spatial layout recovery from intensity estimation. The proposed LiftXR framework interleaves a layout lifter (generating 3D anatomical layout from two X-ray views), an intensity renderer (reconstructing CT volume guided by that layout), and an anatomical parser (refining the layout from the reconstruction), creating a feedback loop between spatial structure and intensity calibration. This is architecturally distinct from end-to-end encoder-decoder approaches and directly addresses the depth-ambiguity problem in limited-view reconstruction. The downstream segmentation improvement on external data suggests the reconstructed volumes carry genuine anatomical fidelity rather than plausible-looking artifacts.

GSToken: Geometry-Structured Gaussian Tokens for Compact 3D Medical Image Representation introduces a tokenization scheme where each token carries a learned 3D center, anisotropic scale, and orientation -- effectively embedding explicit geometric support into the representation at negligible parameter cost. Applied to multi-modal brain tumor segmentation, this is the first use of Gaussian tokens in that domain. The authors also propose a frozen-token utility evaluation protocol that isolates tokenizer information retention from decoder capacity, addressing a real methodological gap in how tokenizers are benchmarked. The consistent outperformance of capacity-matched adaptive baselines under controlled probing conditions suggests that encoding spatial geometry within tokens is a meaningful design principle for compact volumetric representations, not just an incremental architectural tweak.

Until the next announcement!

P.S. In the most recent submission window: physics.med-ph: 4 papers cs.CV: 98 papers