Hi all,

Here is what's been going on in medical physics and related areas over the last 24 hours.

Here's what you folks might find interesting today.

TOP PICKS

RADIOTHERAPY AND HADRON THERAPY / INDUSTRY-RELEVANT AND TRANSLATIONAL WORK

Closed-loop coupling of personalised and foundation models for real-time treatment guidance with MRI

This is the clearest fit for real-time image guidance in radiotherapy today. The key idea is a *closed-loop* system that combines a patient-specific temporal predictor with a streaming foundation-model segmenter, so the predictor is not just extrapolating motion blindly but is continuously corrected by anatomy-aware supervision. That is a meaningful architectural contribution because it explicitly addresses the failure mode of personalised motion models drifting away from anatomically plausible states during treatment. The paper also ties the method to a clinically relevant endpoint, showing reduced *dosimetric error* at a 400 ms prediction horizon rather than stopping at image similarity metrics. Validation on both a digital phantom and intrafraction MRI from MRI-guided radiotherapy patients makes this feel much closer to deployment than a typical proof-of-concept. If you care about adaptive guidance, latency compensation, and robust motion-aware treatment delivery, this looks very worth reading.

IMAGING PHYSICS AND RECONSTRUCTION

High-dimensional Embedding Prior for Noisy K-space Domain MRIReconstruction

This paper is relevant because it tackles MRI reconstruction directly in *noisy k-space*, where many diffusion-based inverse methods still struggle. The novelty is not just "use diffusion for MRI", but a representation-lifting framework that moves reconstruction into a higher-dimensional embedding space so existing diffusion solvers can operate with a more expressive prior without changing their optimization machinery. That makes it potentially more transferable than a solver-specific redesign, especially if you are interested in learned inverse operators or modular ML-physics hybrids. The authors also claim the biggest gains appear in high-noise regimes and connect that to a theoretical error-propagation analysis, which is more informative than a pure benchmark paper. Because the method is framed as model-agnostic and k-space-native, it may be useful for practical acquisition settings where corruption and undersampling coexist. This looks like a solid reconstruction paper to inspect for whether the embedding trick is genuinely new or mainly a reparameterization.

AnF-DiffPET: Anatomy- and Frequency-Guided Diffusion for PET/CT Denoising

This one sits squarely in ML-for-imaging-physics, with a diffusion framework for low-dose PET denoising that tries to preserve both anatomy and uptake fidelity. What stands out is the combination of CT-conditioned anatomical guidance with explicit *frequency-domain* constraints, rather than relying on a generic image-space denoiser. The proposed components, especially Anatomical-Frequency Guidance and the frequency-contrastive hard-mining objective, suggest the authors are trying to control where the model recovers detail and where it may hallucinate uptake structure. That matters for PET because denoising quality is only useful if quantitative fidelity survives, not just visual smoothness. The abstract claims improvements across four PET/CT datasets and emphasizes anatomical consistency and quantitative fidelity, which are the right axes for translational relevance. If you are tracking diffusion models that are more tightly aligned with modality-specific physics and clinical constraints, this is a good candidate.

NOVEL ARCHITECTURES AND OBJECTIVES FOR IMAGING

Mirror-Fusion Attention for Reflection-Aware Self-Supervised Representation Learning

This is the most conceptually interesting adjacent-methods paper for imaging representation learning. Instead of enforcing strict flip invariance, the authors argue that bilateral medical data often contain meaningful left-right correspondences that should be modeled softly rather than erased, and they implement that with a lightweight Mirror-Fusion Attention module inside a ViT SSL pipeline. That is relevant because it introduces a geometry-aware prior that is not just another standard augmentation recipe, and it appears to improve not only downstream accuracy but also *calibration* and robustness to reflection perturbations. The method is demonstrated on CheXpert and BraTS, so the medical-imaging angle is not incidental. While this is not a reconstruction or radiotherapy paper, the underlying idea of encoding approximate symmetry without suppressing asymmetry could transfer to segmentation, anomaly detection, or even anatomy-aware foundation models. I would treat it as a useful methods paper with possible downstream relevance rather than a must-read clinical paper.

Back tomorrow with more!

P.S. In the most recent submission window: physics.med-ph: 1 papers cs.CV: 100 papers