Hi 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.

TOP PICKS

Novel Architectures and Objectives for Imaging

[Anatomy-Informed Neural Networks: Encoding Anatomic Priors in Loss and Architecture, with an SE(3) Formulation of Guidewire-Induced Aortoiliac Deformation]

This paper introduces Anatomy-Informed Neural Networks (AINN), drawing a direct analogy to physics-informed neural networks but for anatomical constraints rather than physical PDEs. Soft priors enter as penalty terms in the loss (e.g., a branching penalty that flags a renal transplant artery off the iliac as unexpected rather than impossible), while hard priors are baked into the architecture itself, making anatomically invalid predictions structurally unrepresentable. The method lifts vessel centerlines and wire paths into the Lie group SE(3), couples a Cosserat-rod wire model to a tortuosity-modulated vessel through a unilateral lumen-contact inequality, and uses a Wasserstein-2 optimal transport loss between predicted C-arm projections and observed angiograms to enable 2D-to-3D learning. The prediction is a constrained minimizer of coupled elastic energy with contact forces as Lagrange multipliers, which is a principled variational formulation rather than an empirical loss hack. While the current submission verifies the mechanics solver only against its own optimality conditions and predictions are not yet mesh-converged, the framework is a compelling blueprint for encoding domain-specific geometric and physical structure into network design. This matters for endovascular navigation and any imaging setting where anatomical validity must be guaranteed by construction.

[LHMCF-Net: A Learned Hyperbolic Mean Curvature Flow Network for Medical Images Segmentation]

LHMCF-Net replaces the standard first-order parabolic contour evolution (as in Chan-Vese) with a second-order hyperbolic PDE governing mean curvature flow, where a velocity field provides inertia and momentum to the evolving interface. The hyperbolic formulation lets contours escape noise-induced local minima and propagate coherently through low-contrast or ambiguous regions where parabolic flows typically get stuck. The PDE solver is unfolded into a deep network where each stage corresponds to a physically interpretable dynamical update, inheriting the PDE's stability and geometric consistency while supporting end-to-end optimization. Results on three public medical segmentation datasets show strong gains precisely in the challenging low-contrast and unclear-boundary scenarios where standard architectures struggle. This is a genuinely novel architectural primitive rooted in differential geometry rather than another Transformer or U-Net variant, and the deep-unrolling-of-a-hyperbolic-PDE pattern could transfer to other geometric evolution tasks in imaging.

Imaging Physics and Reconstruction

[Consistency Models for Fast MRI Reconstruction Using Regularization by Denoising]

CM-RED integrates pretrained consistency models into the regularization-by-denoising (RED) framework for accelerated MRI reconstruction, achieving high-quality results with only 4 network function evaluations compared to the dozens typically required by diffusion-based methods. The method builds on accelerated proximal gradient RED (RED-APG) and adds controlled noise injection during update steps to enhance generative diversity and speed convergence. Experiments across fastMRI knee and brain datasets show consistent outperformance of both diffusion- and consistency-model-based baselines across multiple anatomies, contrast weights, acceleration factors, and undersampling patterns. The approach also exhibits strong robustness to hyperparameter variations, which is practically important for deployment across varying acquisition protocols. By cutting the inference budget from dozens of NFEs to four while maintaining generative-prior quality, this work brings diffusion-style reconstruction closer to clinically feasible reconstruction times. The code and pretrained models are publicly available, which should facilitate reproducible comparison.

Radiotherapy and Translational Work

An integrated diffusion-weighted imaging processing and interpretation platform for MR-guided radiotherapy

This work presents an end-to-end platform that takes raw MR-Linac DWI acquisitions through deep-learning-based distortion correction, denoising, and IVIM/ADC fitting, then couples the quantitative output to a retrieval-augmented generation (RAG) agent that produces traceable clinical interpretations. The RAG agent reasons over a two-layer knowledge base of curated publications with line-level citation tracing, delegates arithmetic to deterministic tools, and was independently rated by a physicist and physician across 54 ratings with a pooled mean of 4.65 out of 5. This directly addresses the MR-Linac workflow, where DWI is acquired at every treatment fraction but converting noisy low-SNR data into clinical decisions remains a bottleneck. The emphasis on traceable LLM reasoning with source-level citations is a notable design choice for regulatory-relevant deployment in radiation oncology, where unverifiable model outputs are a barrier to clinical adoption. The platform also highlights the safeguards needed to verify LLM-generated reasoning, which is a translational concern that extends well beyond this specific application.

Until the next announcement!

P.S. In the most recent submission window: physics.med-ph: 2 papers cs.CV: 98 papers