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 / IMAGING PHYSICS AND RECONSTRUCTION

Self-Supervised Implicit CEST Reconstruction via Physics-Informed Lorentz Encoding

This one is especially aligned with interests in **physics-informed reconstruction** rather than generic image-to-image learning. The core idea is to replace standard positional encodings in an implicit neural representation with **Lorentz Encoding**, where sparse CEST sampling coordinates are projected into a latent space constrained by mixtures of Lorentzian spectral profiles. That is a meaningful architectural change because the inductive bias comes directly from the expected spectroscopy physics, not just from a more expressive network. The method is also **self-supervised**, which matters for ill-posed MRI settings where dense ground truth is hard to obtain. Beyond reconstruction metrics, the paper claims the learned latent trajectory is geometrically ordered and preserves downstream quantitative metabolite mapping for APT, NOE, and MT, which makes it more relevant than a pure PSNR paper. For anyone tracking new reconstruction objectives tied to acquisition physics, this looks like one of the strongest papers in today's batch.

PhyMRI-SR: Toward Physics-Aware MRI Image Super-Resolution

I think this is notable because it reframes MRI super-resolution as a **physics-aware reconstruction problem** where resolution is not fixed, but coupled to SNR through acquisition trade-offs. The unusual part is the use of **2D Gaussian Splatting** as a coordinate-based, resolution-agnostic rendering framework for MRI, which is a more original design choice than another UNet or ViT variant. On top of that, the model predicts intrinsic tissue parameters like **proton density** and **effective relaxation rate R2**, then synthesizes image intensities through governing equations, so the contrast formation is constrained by MRI physics rather than learned purely empirically. The anatomical and imaging-system priors also make the representation more hardware- and tissue-aware, which could be useful for translational reconstruction pipelines. The meta-learning setup for simulated-to-real adaptation is less novel than the physics layer, but it does help with the paired-data bottleneck. Overall, this feels relevant both for **novel imaging architectures** and for **ML-physics hybrids in MRI reconstruction**.

RADIOTHERAPY AND HADRON THERAPY / INDUSTRY-RELEVANT AND TRANSLATIONAL WORK

WING: A Window-Prior-Based Generative Network with Gated Inception for Cross-Modality CT Synthesis

This is the most directly relevant radiotherapy paper today because it targets **MRI-to-CT and CBCT-to-CT synthesis** for adaptive workflows. What stands out is the reformulation of CT synthesis into **multiple windowed targets** instead of direct regression over the full Hounsfield range, motivated by the idea that clinically meaningful structures are better separated in window space than in the raw long-tailed intensity distribution. That is a more thoughtful objective design than standard voxelwise synthesis, and it is paired with a **Gated Inception Generator** plus a **Fuse-and-Refine Transformer** to merge the window-specific predictions back into full-range CT. The paper is still a synthesis model rather than a full treatment-planning system, but the application is clearly tied to **adaptive radiotherapy** and reducing extra radiation exposure. I also like that it claims a single compact model across MRI-to-CT and CBCT-to-CT and across anatomies, which is the kind of practical flexibility that matters for deployment. If you are

watching for translational work that could plug into planning or dose workflows, this is probably the one to read first.

UNCERTAINTY QUANTIFICATION AND LIKELIHOOD-FREE INFERENCE

TMF-RSE: Tri-Modal Fusion with Regional Semantics and Evidential Uncertainty for Lung Severity Scoring

This is more application-heavy than the reconstruction papers above, but I am flagging it because it explicitly builds in **evidential uncertainty** rather than treating uncertainty as an afterthought. The model combines three information sources: image appearance, segmentation-derived structural cues, and semantic features from a vision-language model, then uses evidential regression to output both severity scores and uncertainty estimates. That multimodal design is not entirely new, but the uncertainty formulation could be useful if you are interested in **calibrated prediction under noisy or ambiguous imaging conditions**. The regional semantics angle also makes the fusion more structured than a simple concatenation baseline. It is not as methodologically deep as simulation-based inference or conformal methods, so I would rank it below the top two papers, but it is still one of the more relevant uncertainty-aware medical imaging papers in today's list. Worth a look if you want ideas for uncertainty-aware clinical scoring pipelines.

Back tomorrow with more!

P.S. In the most recent submission window: physics.med-ph: 4 papers cs.CV: 100 papers