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

## UNCERTAINTY QUANTIFICATION AND LIKELIHOOD-FREE INFERENCE

### [ConRad: Efficient Conformal Prediction for Radiomics]

Cheung, Veeraraghavan, and Balakrishnan tackle a very practical uncertainty problem that sits downstream of segmentation: how to put valid intervals on *radiomic features* computed from predicted masks. The key novelty is that their conformal predictor is not a generic black-box wrapper, but an *adaptive* one that conditions on test-time covariates such as mask geometry, image appearance, predicted radiomics, and especially boundary uncertainty. That matters because standard conformal intervals can be valid yet too wide to be useful, whereas this method tries to preserve finite-sample coverage while recovering efficiency from segmentation-specific information. The ablation result is also notable: boundary uncertainty contributes the most to interval tightening, which is exactly the kind of mechanistic insight that makes the method more than a benchmark tweak. For anyone interested in calibrated uncertainty in imaging pipelines, this is one of the clearest translational papers in today's list because it connects segmentation reliability directly to downstream quantitative biomarkers.

## IMAGING PHYSICS AND RECONSTRUCTION

### [Beyond Thermal Imaging: Inferring Thermophysical Properties from Time-Resolved Thermal Observations]

This is not a medical imaging paper per se, but it is a strong adjacent-methods paper with clear transfer potential because it frames inverse imaging as *physics-constrained parameter inference* rather than pure appearance reconstruction. The proposed ThermoField represents geometry and spatially varying thermophysical properties as neural fields, then ties them together with a differentiable heat-transfer simulator and temporal thermal observations. What makes it relevant is the combination of learned scene representation with a governing PDE solver, which is very much in the spirit of forward-model-aware reconstruction and parameter estimation in CT, MRI, PET, or optical imaging. The paper appears to move beyond simply reconstructing what is visible and instead estimates latent physical quantities that govern future dynamics under unseen conditions. That kind of learned inverse-physics formulation is exactly the sort of idea that could transfer to quantitative medical imaging, where one wants tissue parameters or acquisition-physics variables rather than just prettier images. I would flag this as a useful methods paper to mine for architectural and objective-design ideas.

## IMAGING PHYSICS AND RECONSTRUCTION

### [Track2Map: Online Deformable SLAM with Motion-Aware Pose Optimization in Robotic Surgery]

Track2Map is one of the more technically substantive surgical vision papers here because it addresses a real bottleneck in deformable 3D reconstruction: dependence on accurate camera trajectory priors. The method jointly optimizes camera pose and a deformable 3D Gaussian Splatting scene representation *online* from surgical video, which makes it more like a true SLAM system than an offline reconstruction pipeline. The most interesting part is the motion-aware stabilization strategy: dense 2D tracks are used both to initialize deformation and to separate tissue motion from camera motion, including detection of static-camera periods to reduce drift. That is relevant to image-guided intervention and intraoperative imaging because robustness to missing or noisy priors is often what determines whether a method survives contact with the clinic. While it is not centered on CT or MRI, it does fit the broader interest in geometry-aware imaging systems and physically meaningful scene modeling under deformation.

## IMAGING PHYSICS AND RECONSTRUCTION

**Progression as Latent Drift: Generative Forecasting of Slow-Evolving Pathologies**
This paper is aimed at longitudinal brain MRI forecasting, but the reason it stands out is that it does more than apply a generic diffusion or transformer model to disease progression. The authors explicitly analyze why standard generative sequence models fail in the low-signal longitudinal regime, naming two failure modes: *identity collapse* and the *continuous interpolation trap*. Their proposed fix is to predict progression in a compressed semantic latent space rather than at full image resolution, then discretize the learned change representation with finite scalar quantization to suppress nuisance fluctuations while preserving consistent structural drift. That is a meaningful modeling choice for subtle pathology evolution, where most voxels are stable and the clinically important signal is weak and localized. If the claims hold up, this could be useful beyond neurodegeneration, since many medical imaging forecasting problems have the same low-signal, high-nuisance structure.

Back tomorrow with more!

P.S. In the most recent submission window: physics.med-ph: 2 papers cs.CV: 84 papers