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

A Dual-domain Refinement Network with FBP-based Jacobian Learning for Sparse-view Dual-Energy CT Material Decomposition

This one stands out because it does more than apply a generic network to DECT decomposition: it explicitly casts sparse-view multi-material decomposition as a nonlinear least-squares problem and builds the iteration around a learnable approximation to the adjoint Jacobian. The key novelty is the **FBP-based Jacobian approximation module**, which ties the backward step to the imaging physics instead of relying on a purely black-box unrolled update. I also like that the regularization is genuinely **dual-domain**, combining image-space geometric feature extraction with Fourier-domain suppression of sparse-view noise and streak artifacts. That makes it especially relevant if you care about reconstruction methods that respect the forward model while still capturing global structure. If the method works as advertised, it could be a useful template for other nonlinear inverse problems in CT and spectral imaging where the Jacobian matters but is awkward to compute exactly. Among today's papers, this is the clearest fit to physics-aware reconstruction and novel imaging objectives.

DrivenMorph: Bridging Attention Mechanism and Variational Image Registration via Difference Modeling

This paper is interesting because it tries to recover some of the interpretability that deep registration methods often lose by reintroducing a **Demons-like driving force** inside a learned framework. Rather than treating attention as a generic feature mixer, it uses latent-space differences to define an explicit force term that then feeds a **neural Demons layer** for smooth deformation generation. That separation between difference modeling and deformation is a nice design choice: it gives a more inspectable registration pipeline and a clearer physical story for why the deformation field should move the way it does. For researchers interested in geometric or physically motivated registration architectures, this is more compelling than another end-to-end registration U-Net with slightly better Dice. The abstract also claims that the learned driving force aligns with actual deformation patterns, which, if borne out, would make this a useful bridge between classical variational registration and modern learned models. It feels like a method paper with a real inductive-bias contribution rather than just a benchmark entry.

UNCERTAINTY QUANTIFICATION AND LIKELIHOOD-FREE INFERENCE

Set-Inclusive Uncertainty Modeling for Robust Brain Tumor Segmentation

This is one of the stronger uncertainty papers today because it targets a very practical failure mode in multimodal MRI: missing modalities at inference. The core idea is to represent partial-input features as **Gaussian distributions** whose variance is explicitly tied to the information gap relative to the full-modality representation, instead of pretending incomplete evidence is deterministic. What makes it more than a standard probabilistic encoder is the **set-inclusive ordering constraint** over modality subsets, which enforces a hierarchy of uncertainty consistent with how much information is missing. That is a neat structural prior, and it should matter for reliability in real deployment where modality availability is uneven across sites and patients. For anyone interested in calibrated uncertainty in segmentation rather than just mean performance, this is a relevant contribution. It is also nice to see uncertainty treated as a property induced by missing evidence, not just as an afterthought via test-time ensembling.

Beyond Point Estimates for Glaucoma Visual Field Forecasting with Diffusion Models

I think this is worth attention because it reframes visual field progression as a **distributional forecasting** problem rather than a single best-guess regression task. The use of conditioned diffusion models is not novel by itself, but here the important point is that the model produces **plausible future VF distributions** under irregular follow-up intervals and then evaluates calibration on clinically meaningful summary measures. That makes it relevant to uncertainty-aware prediction in longitudinal medical data, especially where heterogeneity and measurement noise are intrinsic. The abstract also says the diffusion model remains competitive when collapsed to a point estimate, which is important because it suggests the uncertainty modeling is not coming at the expense of predictive utility. This is less about imaging physics than the CT paper above, but it is a solid example of moving from deterministic outputs to clinically interpretable uncertainty. If you are tracking how generative models can support risk-aware clinical forecasting, this is one of the better examples in today's list.

RADIOTHERAPY AND HADRON THERAPY / TRANSFERABLE METHODS FROM ADJACENT FIELDS

Compact deep learning pipeline for particle track reconstruction in the pCT detector system

This is probably the most relevant hadron-therapy-adjacent paper today. It tackles a real bottleneck in **proton CT** by using compact networks for two linked tasks: predicting the next hit position to guide track matching, and estimating incoming proton kinetic energy for downstream reconstruction. I would not call the architecture itself highly novel, since it is based on lightweight MLPs, but the way it injects a **physically informed prior** into the matching stage is practical and directly tied to pCT workflow constraints. The emphasis on excluding ambiguous tracks and meeting clinical time limits also makes it feel more translational than many detector-ML papers. For researchers interested in proton imaging, range-related uncertainty, or ML components that could slot into clinically viable pCT systems, this is worth a look. It is not as conceptually strong as the DECT Jacobian paper, but it is one of the clearer application-facing contributions in the hadron space today.

Back tomorrow with more!

P.S. In the most recent submission window: physics.med-ph: 8 papers cs.CV: 100 papers