Hello folks,

Here is what's been going on in medical physics and related areas the last 24 hours. Here's what you might find interesting today.

Imaging Physics and Reconstruction

Flow Matching-Based PET Image Reconstruction (Hashimoto, Huang, Yokota, Gong)

This paper introduces flow matching as a generative prior for **PET image reconstruction**, moving beyond diffusion-based approaches that require many reverse sampling steps with interleaved data-consistency updates. The key insight is that flow matching can directly estimate clean images from intermediate states, decoupling data-consistency refinement from flow propagation. The authors build PET-FlowDPS by incorporating **Poisson likelihood guidance** with an EM-based preconditioner, and then formulate a model-based reconstruction method that interprets the flow-based prior, PET data refinement, and stochastic propagation within an **approximate Bayesian framework**. Results on [18F]FDG brain PET show improved bias-variance trade-offs across dose levels. This matters because it demonstrates a principled alternative to diffusion-based reconstruction with a cleaner separation of physics-based data fidelity from learned priors, which could simplify the design space for quantitative PET reconstruction pipelines.

Novel Architectures and Objectives for Imaging

Coupled Optimal Transport with Landmark Constraints (Gu, Sun, Xu)

This work proposes a **coupled optimal transport framework** that jointly optimizes a transport plan and a deformation field guided by sparse annotated landmarks, connected through a mutual-consistency constraint. The framework provides a principled bridge between **landmark-based registration** and **transport-based distribution matching**, enabling recovery of geometrically meaningful transport maps from sparse supervision. The authors establish well-definedness in a general variational setting and provide convergence analysis for their finite-element numerical scheme. While demonstrated on shape matching, the formulation is directly transferable to **image registration** problems where anatomical landmarks are available but dense correspondence is ill-posed, making it relevant to anyone working on physics-informed or geometry-aware registration in medical imaging.

A Plug-in Interpretation of Conditioning in Score-Based Diffusion Models (Chen, Ghosh, Deveney, Budd, Namboodiri)

The authors propose a conditioning mechanism for diffusion models based on **multi-speed joint diffusion** of the target and condition, learning an unconditional joint score network and enforcing conditioning at inference via a **plug-in correction term**. This separates the conditioning contribution from learned unconditional dynamics, offering a transparent view of how conditions steer generation. They derive explicit conditional reverse-time SDEs and approximate probability-flow ODEs, and introduce a **log-Fokker-Planck residual regularization** to reduce ODE-SDE discrepancy. This is relevant to **image reconstruction and synthesis** because the plug-in formulation could enable more principled and modular conditioning on imaging physics (e.g., forward operators, measurement data) within diffusion-based reconstruction frameworks, and the theoretical grounding makes the conditioning mechanism directly comparable across samplers.

Uncertainty Quantification and Likelihood-Free Inference

Does Marginal Coverage Guarantee Class-Conditional Safety for Zero-Shot VLMs Under Shift? (Sharma, Dutta)

This paper audits **split-conformal prediction** as an abstention layer for zero-shot vision-language models under deployment shift, revealing that marginal coverage can remain high while **class-conditional tail coverage collapses**. On ImageNet-Sketch, worst-class coverage falls to approximately zero despite marginal coverage of ~0.86, with 10-12% of classes below a finite-sample null floor. The authors show that source-side Mondrian calibration does not transfer, and that clustered conformal and Conf-OT improve marginal metrics without recovering the worst-class tail. This matters for **calibrated uncertainty in diagnostic prediction** because it demonstrates that marginal conformal guarantees are insufficient safety mechanisms when deployed in clinical settings with class imbalance or distribution shift, motivating the need for class-conditional or group-conditional calibration strategies in medical imaging pipelines.

Back tomorrow with more!

P.S. In the most recent submission window: physics.med-ph: 2 papers cs.CV: 80 papers