Hey 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.

---

**Imaging Physics and Reconstruction**

PAGS: Autofocusing Photoacoustic Tomography via Speed-of-Sound-Adaptive Gaussian Splatting introduces a genuinely novel architectural primitive for **photoacoustic tomography** reconstruction. Instead of optimizing dense physical medium models for speed-of-sound (SoS) heterogeneity, PAGS represents the initial pressure field with sparse Gaussian photoacoustic sources and replaces explicit medium recovery with a compact anisotropic path-averaged SoS field parameterized by spherical harmonic probes. The analytic Gaussian acoustic projection maps source representations to transducer signals efficiently, enabling closed-loop signal-domain optimization without calibrated SoS priors. This is a **physics-informed architecture** that is not reducible to standard CNN/Transformer variants, and it directly addresses a fundamental ill-posedness in acoustic reconstruction. The computational benefits from the analytic projection and robustness to sparse-view sampling make it practically relevant beyond a benchmark exercise.

FlowMoDL: Model-Based Deep Learning with Conjugate-Gradient Data Consistency for Highly Accelerated 4D Flow MRI Reconstruction presents an unrolled network for **4D flow MRI** that alternates a learned (3+1)D spatiotemporal denoiser with conjugate-gradient data-consistency updates based on the SENSE forward model. The dual-pathway conditioning scheme that adapts both denoiser features and data-consistency weighting is a thoughtful architectural choice, allowing a single model to handle 10x to 50x acceleration. The **velocity-aware composite loss** with curriculum scheduling explicitly penalizes phase-derived angular errors, which matters because standard magnitude losses miss the physics of flow encoding. Under equivalent gradient-step budgets, competing flow-specific networks degrade while FlowMoDL converges robustly, suggesting the CG data-consistency block is doing real work rather than acting as a learned post-processor.

Nanophotonic control of spatial information in scintillation detectors takes a fundamentally different approach to the **dose-resolution tradeoff in CT**. Rather than physically confining scintillation light with pixelated structures, the authors integrate a metalens directly onto a bulk scintillator to use nanophotonic wavefront control that preferentially transfers high-spatial-frequency information to the detector. Their experimentally validated model predicts a fivefold dose reduction and 25-fold increase in resolution bandwidth relative to state-of-the-art pixelated scintillators in a CT-relevant geometry. This is **imaging physics** at the detector level, separating efficient X-ray absorption from optical image formation, and it could have translational impact on commercial CT systems if the fabrication scales.

---

**Novel Architectures and Objectives**

THA-Flow Generative Model: Prosthesis Geometry Prediction from Preoperative CT reframes total hip arthroplasty planning as a **one-to-many conditional generation** problem rather than a deterministic regression, which is a conceptually important shift. The model uses conditional flow-matching in a compressed latent space (via AutoencoderKL) with a 3D UNet learning a rectified flow from Gaussian noise to prosthesis geometry under spatial bone conditioning. Repeated sampling preserves component position and alignment while allowing local geometric variation, capturing the genuine clinical ambiguity in surgical planning. While the application is specific, the **flow-matching architecture** for 3D shape generation conditioned on volumetric imaging is transferable to other planning or synthesis tasks in medical imaging.

---

Back tomorrow with more!

P.S. In the most recent submission window: physics.med-ph: 4 papers cs.CV: 98 papers