Hey there,

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:

SPARC: Slice-to-volume Pipeline for Automated Reconstruction of gated 3D+time fetal Cardiac MRI This paper presents a physics-informed slice-to-volume reconstruction (SVR) pipeline for fetal cardiac MRI, combining physics-based reconstruction with deep learning for segmentation and anatomical reorientation. It tackles the challenging problem of reconstructing 3D+time cine volumes from 2D+time slices under unpredictable fetal motion. The method achieves a tenfold reduction in reconstruction time compared to existing frame-wise approaches while improving quality. The pipeline is deployed in a clinical setting and made publicly available as a Docker container, highlighting its **Industry-Relevant and Translational Work**. This is a strong example of integrating **Imaging Physics and Reconstruction** with learning-based components for a clinically deployable tool.

Angle-domain null subtraction imaging from beamformed plane-wave data This work introduces an angle-domain analogue of null subtraction imaging (NSI) for ultrasound, operating on per-angle images from plane-wave compounding rather than requiring channel data. It provides a matrix formulation to distinguish receive-domain and angle-domain NSI, showing that the angle-domain approach can be implemented as a simple post-processing step. The method effectively reduces apparent lateral width while maintaining computation time close to standard delay-and-sum imaging. This is a neat contribution to **Imaging Physics and Reconstruction**, offering a practical and compatible enhancement for ultrasound beamforming without the need for flexible beamformers or channel data access.

EVADE: Evidence-Verified Agentic Diagnosis with Escape Addressing the critical issue of overconfidence in medical vision-language models (VLMs), this paper introduces EVADE, a non-training method to enhance deployment safety. The method uses a consistency gate across different image views (zoomed versus whole image) to verify predictions and allows for calibrated abstention when views disagree. It directly tackles **Uncertainty Quantification and Likelihood-Free Inference** by reducing expected calibration error without sacrificing accuracy, unlike chain-of-thought or self-consistency methods. This is highly relevant for building reliable diagnostic AI systems, ensuring that models only commit when evidence from multiple views aligns.

A Critical Synthesis of Uncertainty Quantification and Foundation Models for Semantic Segmentation This paper provides a systematic evaluation of uncertainty quantification (UQ) methods integrated with a segmentation foundation model (SAM2). It benchmarks Monte Carlo Dropout, Deep Sub-Ensemble, Test-Time Augmentation, and Evidential Deep Learning across multiple datasets and out-of-domain settings. The work is highly relevant to **Uncertainty Quantification and Likelihood-Free Inference**, revealing clear trade-offs between predictive performance, reliability, and computational cost. It highlights the current limitations of uncertainty-aware foundation models and points to the need for joint optimization of accuracy and robustness for real-world deployment.

Back tomorrow with more!

P.S. In the most recent submission window: physics.med-ph: 3 papers cs.CV: 90 papers