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.

Top Picks

Imaging Physics and Reconstruction

Eddeep: a deep-learning framework for fast eddy-current distortion correction in diffusion MRI tackles a long-standing bottleneck in diffusion MRI processing by replacing the iterative FSL Eddy pipeline with a single-forward-pass deep learning approach. The key novelty is a two-stage decomposition: first, a supervised image-translation network normalizes contrast differences between diffusion-weighted and b=0 volumes, and second, an unsupervised registration network estimates both eddy-current distortion and head motion under a **physics-constrained quadratic distortion model**. This matters because the physics constraint replaces generic registration with a model grounded in the known functional form of eddy-current artifacts, and the authors show comparable correction quality to FSL Eddy across multiple metrics while drastically cutting inference time. The out-of-domain evaluation on Memodyn data (trained on UK Biobank) is a useful generalization check.

Amplitude-based echo-mode speed-of-sound tomography proposes a fundamentally different signal extraction strategy for ultrasound speed-of-sound (SoS) tomography. Instead of the conventional differential time-of-flight (dToF) approach, which reduces angle-dependent echo variations to a single phase-shift parameter, the authors invert a forward model that relates full aberration profiles to echo amplitudes across all transmit-receive angle pairs. This **amplitude-based inversion** achieves sub-beamwidth resolution and substantially reduces artifacts in media with short-scale SoS variations, as demonstrated in simulations, phantoms, and preliminary in vivo data. The work is a clean example of rethinking the measurement model in ultrasound tomography rather than applying ML to the existing pipeline, and the physics phantom validation lends credibility to the approach.

Novel Architectures and Objectives for Imaging

Where Physics Meets Privacy: Federated PINNs for Privacy-Preserving Brain Tumor Biomechanical Modeling combines **federated learning with a physics-informed loss** built on linear elasticity equations for brain tumor biomechanical modeling. Each simulated clinical site trains a local network on patient-specific MRI data using the PDE-constrained loss, sharing only weights via FedAvg. The interesting claim is that the federated model (91.4% accuracy) slightly outperforms a pooled-data baseline (90.0%), suggesting that the physics regularization acts as an effective inductive bias that complements rather than conflicts with distributed training. The smooth, divergence-free displacement fields produced during training are consistent with expected tissue deformation, which is a meaningful sanity check that the physics constraints are doing real work. While the evaluation is on simulated data, the framework demonstrates that **physics-informed objectives and federated protocols can be paired without performance degradation**, which is relevant for institutions that cannot pool patient imaging data.

BATS: Resource-Efficient Volumetric Segmentation with Boundary-Aware Mixed-Resolution Tokens introduces an **architectural primitive for 3D medical image segmentation** that is not a straightforward Transformer or CNN variant. The core idea is a dense boundary predictor that identifies where fine-resolution processing is needed at every scale level, combined with a fine-first context cascade that constructs an input-dependent mixed-resolution token hierarchy. Homogeneous regions are represented coarsely while boundaries, thin structures, and small targets retain fine tokens, yielding over 53% peak GPU memory reduction on KiTS, LiTS, and BraTS relative to MedNeXt-L while staying within 0.37 Dice points on average. The per-level independent boundary prediction is a deliberate design choice to prevent coarse-scale errors from propagating downward and suppressing fine-scale evidence, which is a principled response to a known failure mode in hierarchical architectures. The parent cluster attention mechanism provides cross-scale context without dense multi-scale feature maps, making this a genuinely different computational structure rather than a tuned variant of existing backbones.

Until the next announcement!

P.S. In the most recent submission window: physics.med-ph: 4 papers cs.CV: 91 papers