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.

Novel Architectures and Objectives for Imaging / Imaging Physics and Reconstruction

PRIME-SVR: Physics-infoRmed Implicit Multi-Echo Slice-to-Volume Reconstruction for Fetal T2 mapping introduces an implicit neural representation framework that jointly reconstructs high-resolution fetal brain volumes across multiple echo times. The key novelty is a Bloch equation-derived regularization term that penalizes deviations from expected T2 decay across echoes, with adaptive weighting that strengthens coupling for degraded stacks. A second network estimates slice-specific acquisition degradations, making the approach fully self-supervised. Validated on 39 in vivo acquisitions across two centers, two vendors, and two field strengths (1.5 T and 0.55 T), it achieves 47% sharper reconstructions than state-of-the-art SVR and enables the first INR-based T2 maps at 0.55 T. The physics-informed cross-TE coupling is a genuinely new objective for SVR, not just an architectural tweak, and the acceleration from 15 to 5 minutes while maintaining T2 accuracy is clinically meaningful.

Uncertainty Quantification and Likelihood-Free Inference / Imaging Physics and Reconstruction

SIINR: Structurally Informed Implicit Neural Representations for super-resolution with uncertainty quantification of clinical quality diffusion MRI datasets combines a supervised 3D U-Net prior with a self-supervised INR that fuses the prior and low-resolution data for diffusion MRI super-resolution. The INR provides analytic approximate posterior distributions, enabling principled uncertainty quantification in the reconstructed outputs -- a feature largely absent from standard super-resolution approaches. The framework enforces data consistency and jointly models spatial and angular domains, which matters for downstream microstructure estimation. Validation on clinical cases including multiple sclerosis and brain lesions shows the method can propagate intensity changes and flag uncertain regions, which is critical for clinical trust. The modularity and analytic posterior make this a strong contribution to calibrated uncertainty in image reconstruction.

Uncertainty Quantification and Likelihood-Free Inference / Transferable Methods from Adjacent Fields

Strong Gravitational Lensing Posterior Sampling in Pixel-Space Using Diffusion Models and Recurrent Inference Machines tackles a high-dimensional, nonlinear inverse problem from astrophysics by combining diffusion-based generative modeling with recurrent inference machines (RIMs) to generate joint posterior samples of source brightness and foreground mass distribution. The method produces pixelated posterior samples conditioned on observations, modeling realistic lensing simulations down to the noise level. While the application is gravitational lensing, the methodological core -- amortized posterior sampling via diffusion models coupled with iterative refinement -- is directly transferable to ill-posed inverse problems in medical imaging such as PET reconstruction or MR fingerprinting, where high-dimensional posteriors over latent images are needed. The RIM-diffusion coupling is architecturally novel and not reducible to standard conditional diffusion.

Novel Architectures and Objectives for Imaging / Imaging Physics and Reconstruction

SHFormer: Dynamic Spectral Filtering Convolutional Neural Network and High-pass Kernel Generation Transformer for Adaptive MRI Reconstruction addresses a known limitation of attention-based MRI reconstruction models: their bias toward low-frequency content at the expense of high-frequency detail. The proposed architecture pairs a spectral filtering CNN that captures mode-specific transferable features with a dynamic high-pass kernel generation transformer that explicitly propagates high-frequency information. A neuromodulation mechanism enables cross-domain generalization across heterogeneous MRI modalities without mode-specific retraining, which is a practical bottleneck in clinical deployment. Evaluation covers supervised, self-supervised, and diffusion-based training paradigms, plus open-set generalization, giving a thorough picture of robustness. The high-pass kernel generation is a targeted architectural primitive rather than a generic transformer variant, and the cross-modal transferability results are relevant for multi-center imaging workflows.

Back tomorrow with more!

P.S. In the most recent submission window: physics.med-ph: 4 papers cs.CV: 83 papers