Hello folks,

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.

physics.med-ph

Bridging the gap: Using deep learning to reconstruct noise-reduced super-resolved OCT images from gapped spectra This paper tackles a practical problem in **imaging physics and reconstruction**: combining multiple, affordable, lower-bandwidth optical coherence tomography (OCT) sources often results in spectra with discontinuities. The authors propose a neural network that can reconstruct super-resolved, noise-reduced B-scans from a flexible number of spectra with arbitrary, non-overlapping gaps. The method preserves fine, low-contrast details and edges, and the noise reduction actually increases with the size of the band gap. The approach can also be viewed as a Fourier-domain masked autoencoder for self-supervised image quality enhancement in broadband settings. This could be a major step towards high-quality OCT imaging using spectrally disjoint, low-bandwidth sources.

cs.CV

Evaluating and Calibrating Diffusion Model-derived Uncertainty for Quantitative MRI Mapping This work directly addresses **uncertainty quantification** in the context of **imaging physics and reconstruction**. The authors systematically evaluate uncertainty maps for quantitative MRI (qMRI) derived from multiple inferences of a data-consistent diffusion model. They find that while the diffusion model-derived uncertainty is positively associated with mapping error and useful for selective prediction, it is poorly calibrated for quantitative interval interpretation. To fix this, they introduce a post-hoc calibration procedure combining prediction-value-dependent bias correction with scalar uncertainty scaling, which substantially improves calibration. This is a solid contribution for anyone interested in making deep learning-based qMRI mapping clinically reliable.

M-Net: Integrating Spectral Features and Physical Field Operators into Deep Learning for Medical Image Segmentation This paper introduces a novel architecture that embeds explicit mathematical and physical inductive biases into a U-Net framework, fitting squarely into **novel architectures and objectives for imaging**. M-Net integrates continuous spectral features derived from the condition number of local pixel matrices, alongside physical field operators like divergence and a discrete curl-like boundary irregularity operator. These features are fused with CNN-extracted deep features using a Math-Attention Gate at skip connections. The results demonstrate that these mathematical priors provide complementary information, outperforming baseline U-Net by significant margins on liver, kidney, and brain tumor segmentation tasks.

A Neighborhood Attention Transformer Network for Enhanced 3D Segmentation of the Left Anterior Descending Artery This work is highly relevant to **radiotherapy and hadron therapy**, specifically for automated contouring to enable cardiac dose sparing. The authors propose NA-UNETR, a 3D transformer that uses Neighborhood Attention and Dilated Neighborhood Attention blocks to capture both fine structural details and long-range context for segmenting the Left Anterior Descending (LAD) artery in low-contrast, non-contrast CT. To handle the scarcity of annotated LAD data, the model is pretrained on 1,000 CTA volumes and fine-tuned using LoRA. It also employs a composite Dice-Focal and Hausdorff loss dynamically balanced via homoscedastic uncertainty, showing improved boundary accuracy and centerline stability over standard baselines.

Until the next announcement!

P.S. In the most recent submission window: physics.med-ph: 3 papers cs.CV: 100 papers