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

Physics-Informed Deep Learning Model for Cross-Modality Super-Resolution in Fluorescence Microscopy

This paper tackles a recurring problem in generative cross-modality translation: visually plausible outputs that violate the underlying image-formation physics. The authors embed microscope-specific **point spread function (PSF)** information directly into the GAN training objective for confocal-to-STED super-resolution, making the loss **forward-model-aware** rather than purely empirical. They validate on both simulated and experimentally measured PSFs using paired confocal-STED data of TOM20-labeled mitochondria in human macrophages, and assess results with frequency-domain and distribution-sensitive metrics that probe physics-relevant properties such as spatial-frequency content and SNR behavior. The PSF-guided models improve structural fidelity and reduce local deviations versus non-PSF baselines, particularly in frequency-domain analyses. This matters because it demonstrates a concrete route to making generative microscopy models physically consistent, not just visually convincing, which is essential if such methods are to be trusted for quantitative biological measurement.

Flash EQ-Linear: Accelerating Equivariant Linear Layers via Group-wise Discrete Fourier Transform

Equivariant networks embed geometric symmetries as structural priors and are highly parameter-efficient, but existing implementations unroll structured weights into dense matrices, so compute cost matches a non-equivariant layer. The authors observe that the equivariant linear layer is fundamentally a circular convolution along the group dimension composed with a channel-wise linear transform, and exploit the **Fourier convolution theorem** plus conjugate symmetry of the real DFT to reduce complexity from O(NDC) to O(NDC/T). They provide dedicated CUDA kernels for forward and backward passes in FP32 and FP16, achieving up to 2x forward speedup at the operator level and 1.7x end-to-end speedup for equivariant ViT and Swin variants. This is the first demonstration that equivariant networks can simultaneously dominate non-equivariant counterparts on accuracy, parameters, and inference speed. For medical imaging groups considering equivariant architectures for reconstruction or segmentation under rotational or reflectional symmetry, this removes the practical compute penalty that has limited adoption.

Latent Variable-Mediated Cross-Learning for Few-Shot Acoustic Impedance Imaging

Acoustic impedance imaging is severely ill-posed: the seismic wavelet is unknown, observations are band-limited, and labeled samples are extremely scarce. The authors propose RD-SCL, which integrates a **differentiable, closed-form first-order Tikhonov deconvolution operator** that dynamically estimates the latent wavelet in the frequency domain during training, providing **physics-guided feedback** without auxiliary networks or fixed wavelet priors. A symmetric cross-learning scheme enforces consistency between labeled and unlabeled predictions, exploiting abundant unlabeled traces. The method uses only 56.5k parameters and outperforms supervised and semi-supervised baselines on the SEAM and Marmousi 2 benchmarks. While the application is geophysical, the core methodology, a differentiable regularized inverse operator embedded in the training loop with semi-supervised consistency, is directly transferable to **medical image reconstruction** problems where the forward model is known but parameters must be estimated jointly with the image.

UnDA: Unpaired Domain Alignment for Cross-Modal Knowledge Transfer in Medical Imaging

Cross-modal knowledge distillation in medical imaging often fails when the modality gap is large and source-domain predictions are noisy. The authors introduce **Uncertainty-Weighted Optimal Transport (UCT-OT)**, which dynamically weights feature-level alignment based on prediction confidence, suppressing noisy supervision from uncertain source predictions. An attention-based pooling mechanism extracts semantically structured class tokens in a backbone-agnostic alignment module, and a per-class ProtoNCE objective maintains stable prototype memories for global discriminability across unpaired batches. Evaluations on representative segmentation tasks under strictly unpaired settings show consistent improvements in accuracy and boundary precision. This is relevant because **optimal transport** and **uncertainty quantification** are combined in a principled way to handle the practical reality that paired multi-modal clinical data is rarely available, and the information-geometric grounding of the transport plan makes the alignment more robust than heuristic domain adaptation.

Until the next announcement!

P.S. In the most recent submission window: physics.med-ph: 0 papers cs.CV: 100 papers