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.

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

Exact and Calibrated Diffusion Reconstruction for Digital Breast Tomosynthesis tackles three persistent clinical obstacles in limited-angle DBT reconstruction with conditional diffusion priors: inexact data consistency, unlocalized hallucination, and uncalibrated uncertainty. The key novelty is replacing the standard per-step proximal update in a diffusion sampler with an exact Euclidean projection onto the data-consistent set, computed via a dual system with a one-time Gram matrix factorization that costs only 4.5 ms per step (a 248x speedup) and drives residuals to the double-precision floor. The authors prove this projection is the limit of the proximal step as the penalty parameter goes to zero, provide a no-harm theorem, and show that the ensemble variance is supported entirely on the null space of the forward operator, meaning the mean's error lies in the unmeasured subspace. Isotonic recalibration then brings the ensemble spread to a calibrated error scale (ECE drops from 0.029 to 0.008). This is the first data-consistent, uncertainty-calibrated learned reconstruction for limited-angle DBT, and the solver naturally relaxes to discrepancy-ball and MAP modes for noisy measurements.

Imaging Physics and Reconstruction

Reconstruction-Independent Resolution Limits in Preclinical Cone-Beam Micro-CT: A Closed-Form Analysis derives two fundamental, reconstruction-independent bounds on spatial resolution in cone-beam micro-CT: a photon-noise limit from the SKE/BKE ideal observer with the Rose criterion, and an angular-sampling limit from the Crowther criterion. Both are derived in closed form for a circular feature in a uniform background, including an off-axis fan-beam extension, and combined into a spatially varying effective resolution map. The combined limit yields an analytic scan-design optimum where the maximum worst-case spatial frequency scales as the cube root of the total photon budget per detector pixel, with an explicit rule for splitting that budget between per-view flux and number of views. The framework is illustrated on a representative preclinical scanner with parametric sweeps, and an open-source web calculator implements the bounds for arbitrary geometries. This is a clean theoretical contribution that gives concrete, actionable design guidance grounded in imaging physics rather than empirical tuning.

Novel Architectures and Objectives for Imaging

Medical Image Segmentation based on Deep Active Contour and Mean Curvature Loss Function introduces a geometric prior into deep learning segmentation by proposing the Deep Active Contour and Mean Curvature (DACMC) loss, which augments the Chan-Vese region-based loss with a mean curvature term. The mean curvature is approximated using convolution kernels to keep computation tractable, providing a natural geometric constraint on the segmentation boundary that prior Chan-Vese-based losses lack. The method is validated on liver and spleen datasets and reports new state-of-the-art performance. While the gains are empirical, the conceptual contribution of embedding curvature regularity directly into the loss function is a meaningful step toward physics- and geometry-aware objectives for segmentation, aligning with the broader goal of moving beyond purely pixel-level training.

Uncertainty Quantification and Likelihood-Free Inference

Calibrated Selective Prediction Using Deep Ensembles for ROI-Based Thyroid Nodule Ultrasound Classification Under Dataset Shift presents a calibrated five-member deep ensemble for thyroid nodule malignancy classification with a selective referral policy based on mutual information as an ensemble-disagreement score. The framework uses member-wise vector-scaling calibration and a three-tier triage policy (No-FNA, FNA, radiologist review). On internal cross-validation the ensemble achieves strong discrimination (AUC 0.94) and calibration (ECE 0.009), but the critical finding is the substantial degradation under external dataset shift (AUC drops to 0.79, ECE rises to 0.19), with the frozen policy assigning 83.7% of external cases to radiologist review. This is a valuable, honest translational study: it demonstrates that selective prediction can identify cases unsuitable for automated triage, while also exposing the limited transportability of thresholds and calibration across institutions, underscoring the need for local recalibration before clinical deployment.

Until the next announcement!

P.S. In the most recent submission window: physics.med-ph: 2 papers cs.CV: 100 papers