Hey there,

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

### Uncertainty Quantification and Likelihood-Free Inference

Score-Based Ideal Observer Approximation via Denoising Score Matching for Signal-Known-Exactly Detection Tasks -- This paper reformulates the Bayesian Ideal Observer (IO) test statistic in terms of the score function (gradient of the log probability density), introducing a score-based ideal observer (SIO) that is trained exclusively on signal-absent images via denoising score matching. The key novelty is that once trained, the score model can approximate the IO test statistic for *arbitrary additive signals* without per-image posterior sampling or signal-specific retraining, which is a significant departure from MCMC-based approaches and supervised learning methods that are task-specific. This matters because it bridges **likelihood-free inference** with **imaging physics** detection theory, potentially enabling amortized optimal observer computation across a family of detection tasks. The numerical validation on a stochastic lumpy-background SKE task shows close approximation of IO performance, suggesting the score-function representation captures the essential geometry of the data distribution needed for optimal detection.

It depends: Incorporating correlations for joint aleatoric and epistemic uncertainties of high-dimensional output spaces -- This work tackles a genuinely hard problem in **uncertainty quantification**: jointly modeling aleatoric and epistemic uncertainty in high-dimensional outputs (e.g., medical image segmentation or restoration) by approximating the joint covariance with a low-rank plus diagonal structure. Unlike prior methods that treat the two uncertainty types separately, this approach unifies them into a single second-order distribution that supports downstream sampling and log-likelihood evaluation. The low-rank plus diagonal decomposition captures essential output correlations while remaining computationally tractable, which is critical for **calibrated uncertainty** in clinical imaging pipelines. The stabilization strategies for training and inference are a practical contribution, and the method is validated across image inpainting, colorization, optical flow, and depth estimation, demonstrating broad applicability.

### Imaging Physics and Reconstruction

Effects of Quantum Noise and Source Blurring on Dark-Field Signal Retrieval in X-ray Speckle-Based Imaging -- This is a rigorous **imaging physics** study that systematically characterizes how hardware-induced noise (photon starvation) and source blurring propagate through two mathematically distinct dark-field retrieval algorithms: differential-based intrinsic tracking (LCS) and patch-wise explicit tracking (XST-XSVT). The finding that derivative noise amplification in LCS causes precipitous drops in signal linearity and sensitivity under photon starvation, while XST-XSVT's patch-wise variance calculation inherently suppresses stochastic noise, is a concrete mechanistic insight rather than a marginal benchmark result. This characterization establishes operational boundaries for low-power and low-coherence X-ray systems, directly informing algorithm selection for quantitative dark-field imaging in preclinical and clinical applications. The work bridges **forward-model-aware** analysis with practical **translational** considerations for emerging X-ray modalities.

ZODIAC: Zero-shot Octree-based Diffusion for Anatomical Completion -- This paper proposes a zero-shot shape completion framework that reconstructs the full lumbar spine from partial intraoperative ultrasound without relying on simulated paired training data, which is a notable departure from supervised approaches that assume a predefined occlusion distribution. The method learns a generative diffusion prior over full anatomical shapes in an adaptive octree representation, enabling efficient single-forward-pass completion, and introduces a "blended completion" mechanism that integrates the learned prior with incoming partial geometry at inference time. This is relevant to **novel architectures for imaging** because the octree-based

diffusion prior is a geometric primitive not reducible to standard Transformer/CNN variants, and the zero-shot formulation addresses a fundamental generalization gap when real occlusions diverge from simulated ones. The 22% improvement in HD95 over a fully supervised variant on phantom and volunteer data is a meaningful result for **translational** intraoperative imaging.

---

Until the next announcement!

P.S. In the most recent submission window: physics.med-ph: 4 papers cs.CV: 100 papers