Hey there,

Here is what's been going on in medical physics and related areas over the last 24 hours. A small batch today, but there are some genuinely interesting ideas worth your time. Here's what you folks might find interesting today:

TOP PICKS

Imaging Physics and Reconstruction

Self-Auditing Residual Drifting for Pathology-Preserving Accelerated Knee MRI *Qing Lyu, Jianxu Wang, Mohammad Kawas, Ge Wang, Christopher T. Whitlow*

This paper introduces SA-RDM-DC, which adapts a generative drifting paradigm to accelerated MRI reconstruction in a way that is meaningfully distinct from standard diffusion or score-based approaches. Rather than iterating from noise, it trains a physics-conditioned drift field that moves directly from the zero-filled reconstruction toward a residual correction, with explicit k-space data consistency enforced throughout. What makes this particularly notable is the self-auditing mechanism: dense error maps and slice-level risk scores are produced in the same single inference pass as the reconstruction itself, giving you a built-in reliability signal without any additional compute. The pathology-aware evaluation on fastMRI+ annotations is a welcome departure from SSIM-only benchmarking, showing that lesion-region fidelity and meniscus prediction stability are preserved at acceleration factors up to 12. For anyone thinking about clinically deployable reconstruction pipelines, the combination of subsecond inference, data-consistency enforcement, and case-specific uncertainty flagging is a compelling package.

Radiotherapy and Translational Work

Prediction of Radiotherapy-Induced Hematologic Toxicity in Cervical Cancer with Cohort-Aware Framework *Haotian Feng, Yan Kong, Tiantian Yang, Ke Sheng*

Hematologic toxicity is a clinically significant dose-limiting complication in pelvic radiotherapy, and this paper tackles a problem that is often quietly swept under the rug: radiomic and dosiomic models are highly sensitive to who drew the contours, which kills generalizability across institutions. The cohort-aware neural network (CANN) proposed here learns both shared and cohort-specific representations simultaneously, regularized with contrastive alignment, rather than trying to harmonize away the cohort signal entirely. The key finding is that naive pooling and standard statistical harmonization actually hurt performance, while adversarial alignment degrades it further, suggesting that cohort identity carries genuine signal that should be modeled rather than erased. CANN recovers a test AUC of 0.72 against a dosimetry-only baseline of 0.58, with ablations confirming that both the shared and cohort-specific branches contribute. This is a practically important result for anyone building radiomic toxicity models intended for multi-institutional deployment.

Imaging Physics and Reconstruction

X-Splat: Gaussian Splatting for 3D CBCT Generation from Single Panoramic Radiograph *Tomasz Szczepanski, Szymon Plotka, Michal K. Grzeszczyk, Tomasz Trzciński, Arkadiusz Sitek*

This is a genuinely novel application of Gaussian Splatting to a hard inverse problem in dental imaging: recovering a full 3D CBCT-like volume from a single panoramic X-ray, which is a severely underdetermined problem given that the panoramic acquisition integrates 3D attenuation along curved paths into a 2D image. The key architectural idea is to initialize learnable anisotropic Gaussian primitives along the known X-ray acquisition paths, then refine them in a single feed-forward pass constrained by Beer-Lambert reprojection loss and multi-view

radiographic supervision. This is a physics-informed geometric approach that avoids the oversmoothing typical of NeRF-based methods and the anatomical inconsistencies of GAN-based synthesis. The introduction of segmentation-based geometry-aware metrics for evaluating maxillofacial anatomy reconstruction is also a useful methodological contribution for the field. Recovery of fine structures like the mandibular canal, which prior methods fail on, suggests the geometric scaffold is doing real work here.

Imaging Physics and Reconstruction

Free-breathing Pulmonary Relaxometry at 0.55T *Pavlos Panos, Oliver Bieri, Maurice Pradella, Katrin E. Hostettler, Grzegorz Bauman*

Quantitative lung MRI has long been hampered by the need for breath-holds and external respiratory gating, both of which limit patient compliance and clinical scalability. This work demonstrates a fully automated free-breathing pipeline for simultaneous T1 and T2 mapping of the lung at 0.55T, combining an IR-uf-bSSFP sequence adapted for high-temporal sampling of the transient phase with deep learning-based lung segmentation and multi-contrast registration for motion correction. The low-field 0.55T setting is a deliberate and interesting choice, as it offers longer T2* and reduced susceptibility artifacts in the lung compared to higher field strengths, making quantitative mapping more tractable. The pipeline requires no manual intervention and produces voxel-wise relaxometry maps that agree closely with breath-hold references, and in a patient case successfully distinguished a solid lung mass from surrounding parenchyma by T1 contrast alone. This is a clinically translatable workflow that removes two major barriers to routine quantitative lung MRI.

Until the next announcement!

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