

WEEKLY AI INTELLIGENCE REPORT

Radiology and Medical Imaging

Week 42 | 11 to 17 October 2026

Dr. Sergey Morozov | drsergeymorozov.com

Compiled from publicly available sources (indexed peer-reviewed literature, official press releases, regulatory databases). Does not constitute medical, regulatory, or financial advice.

8

Peer-reviewed papers

4

Industry / FDA items

3

Media highlights

3

Key opinion leaders

GOVERNANCE TAKEAWAY OF THE WEEK

Validate the use, not the tool. This week the same mammography AI swung from modest to large gains depending on how it was wired into the read; adult chest models failed on neonates until they were adapted with 50 local radiographs per site; and the largest prostate MRI AI study to date matched radiologists on average while leaving 28% of sites below an AUC of 0.80. The action: for every deployed model, write a single line naming the integration policy, who sees the output and when it can override, and the patient and reader population your acceptance test actually covered. Treat any use outside that line, a new scenario, a new population, a new site, as a new deployment that needs its own acceptance test and monitoring baseline.

SECTION 0

EXECUTIVE SUMMARY

Key trends, week of 11 to 17 October 2026

[GOVERNANCE] THE INTEGRATION SCENARIO IS THE DEVICE

In a 130,838-mammogram multicenter study, the same breast AI produced materially different reader performance depending on how disagreements were resolved: aggregated reader AUC rose from 83.04% unaided to 88.40% under a union policy for uncertain cases. Acceptance tests have to name and test the deployed decision policy, not the algorithm in isolation.

[EVIDENCE] ADULT MODELS, NEONATAL PATIENTS, A WORKABLE FIX

In Nature Biomedical Engineering, 30 neonatal ICUs in 20 countries each contributed 50 chest radiographs to adapt an adult tube-position model by continual learning, without pooling data: landmark error fell from 4.30 cm (adult-only) to 0.67 cm, better at 29 of 30 sites. Fewer than one in five commercial imaging AI models are cleared for children.

[IMPLEMENTATION] AUDITABLE BY CONSTRUCTION

A two-trust NHS study put open-weight LLMs on extraction only and left the guideline to a deterministic rule engine: postpolypectomy surveillance recommendation accuracy 0.96 and 0.93, cancer alert F1 1.00 at both sites, and the tool improves adherence mainly by catching undersurveillance. A template for keeping LLM pathways auditable.

[WATCHPOINT] PROSTATE AI CLEARS THE BAR ON AVERAGE, NOT EVERYWHERE

SCARLET-1 in Radiology (22,445 prostate MRI exams, 22 countries) found AI comparable to radiologists for clinically significant cancer overall, yet per-site AUC ranged from below 0.80 at 28% of sites to 0.90 or more at 32%. The authors' own conditions for safe use: operating-point selection, local testing and longitudinal performance monitoring.

[OVERSIGHT] VALIDATED FOR WHOM, AND WIRED IN HOW

The week's thread: performance does not travel by itself across sites, populations or integration policies. A mammography aid swung with the decision policy, adult models failed on neonates until adapted, and the biggest prostate MRI AI study left a quarter of sites below AUC 0.80. Write one line per deployed tool naming the decision policy and the population your acceptance test covered; everything outside that line is a new deployment.

SECTION 1

PEER-REVIEWED PAPERS

Source: PubMed | Verified indexing | No preprints | 11 to 17 October 2026

Optimizing AI-assisted decision-making in mammography screening: focusing on uncertainties.

Ahmadinejad N, et al. | *Insights into imaging* | 2026-10-13

Before you wire an AI decision aid into screening mammography, decide and test where it sits in the read. This multicenter study evaluated the same breast AI under five integration scenarios, from unaided reading to AI as the final arbiter of uncertain cases, and the scenario, not the model, determined the gain; the union policy for uncertain cases delivered the best reader performance.

Key metrics: 130,838 mammograms from 35,709 patients with confirmed follow-up outcomes, multiple centers, DDSM as external validation. Aggregated reader AUC 83.04% unaided (95% CI 79.8 to 86.3) vs 88.40% (85.6 to 91.2) in the union scenario, $p < 0.001$; sensitivity 54.8% to 83.9%; specificity 83.6% to 91.1%. Retrospective.

Clinical relevance: The clinically deployed object is not the algorithm but the decision policy around it: the same model produced different accuracy depending on how radiologist and AI disagreement was resolved. For your AI committee this is an acceptance-testing item under EU AI Act Article 14: specify which integration scenario goes live, test that scenario with local readers, and treat any later change of policy, for example from assistant to arbiter, as a new deployment that needs re-acceptance.

PMID 42867728 DOI 10.1186/s13244-026-02411-4

Improving Postpolypectomy Surveillance Decisions Using an Automated Analysis of Colonoscopy and Histopathology Reports: 2-Center Validation Study.

Zecevic A, et al. | *Journal of medical Internet research* | 2026-10-16

If you want large language models inside a guideline pathway, copy this architecture: open-weight LLMs do only the structured extraction from colonoscopy and histopathology reports, and a deterministic rule engine encoding the guideline computes the surveillance interval, so every recommendation is auditable. Developed at one NHS trust and externally validated at a second.

Key metrics: Extraction on 100 procedures (215 specimen pots): specimen type 0.995, location 0.981, diagnosis 0.981, excision completeness 0.949. Surveillance recommendation accuracy 0.96 (95% CI 0.90 to 0.98) at the development trust and 0.93 (0.86 to 0.97) at the external trust; cancer alert F1 1.00 at both. In a 127-patient audit against routine practice, the tool would improve guideline adherence mainly by reducing undersurveillance.

Clinical relevance: Separating probabilistic extraction from deterministic decision logic is what makes the tool auditable and its error modes separable. For your AI committee this changes acceptance testing for any LLM-based pathway tool: demand separate performance evidence for the extraction layer and the decision layer, plus traceable intermediate outputs, in line with the record-keeping and logging expectations of EU AI Act Articles 12 and 19.

PMID 42886590 DOI 10.2196/102789

Curriculum-Aligned Conversational AI to Support Pathology Self-Directed Learning in Preclinical Medical Students: A Crossover Randomized Pilot Study.

Roark K, et al. | *JMIR medical education* | 2026-10-13

A curriculum-aligned custom GPT with structured Socratic prompting beat fully unguided preparation for pathology small-group teaching in a crossover randomized pilot: better self-directed learning documentation and better discussion performance, with an honest null on delayed knowledge retention.

Key metrics: 44 second-year medical students, two cases in crossover design. SDL rubric 10.3 vs 9.4 out of 12 (mean difference 0.9, 95% CI 0.5 to 1.3, $P < .001$, Cohen d 0.71); discussion score 14.9 vs 13.7 out of 16 (difference 1.2, 95% CI 0.7 to 1.7, $P < .001$, d 0.78). Delayed retention on end-of-block exam items not significantly different ($P = .052$). 34 of 44 students reported clearer reasoning. Unregistered single-site pilot.

Clinical relevance: The design detail that matters is alignment: the tool was built on the course's own curriculum rather than used as a general chatbot, and the gains appeared in reasoning and process measures while retention did not differ. A careful small study worth following before extrapolating to residency or continuing education; the authors themselves call for larger multisite trials.

PMID 42873617 DOI 10.2196/102908

Development and External Testing of an Interpretable MRI-Based Machine Learning Model for Differentiating Endometrial Stromal Sarcoma from Cellular Leiomyoma.

Zhou Y, et al. | *Journal of imaging informatics in medicine* | 2026-10-13

Three interpretable MRI features, irregular tumor margin, cystic areas on T2 and mean ADC, carried a support vector machine that separates endometrial stromal sarcoma from cellular leiomyoma, held up in external testing at a fixed threshold, and lifted the junior reader most, in a preoperative call that changes the surgical plan.

Key metrics: 155 women with pathological confirmation, two centers: 111 derivation, 44 external testing. AUC 0.843 (95% CI 0.760 to 0.917) in derivation and 0.835 (0.669 to 0.960) in external testing, threshold fixed in derivation. Exploratory model-assisted reading: junior radiologist specificity 46.4% to 82.1%, accuracy 52.3% to 77.3%. Retrospective.

Clinical relevance: Clinically attractive because the three predictors are features a radiologist can verify on the images rather than a black-box score, and because the external cohort kept the derivation threshold instead of re-optimising it. Prospective multicenter validation is still owed before this guides surgery.

PMID 42870161 DOI 10.1007/s10278-026-02352-w

Interpretable Source-Aware Bone-Muscle Fusion Framework for Multicenter Prediction of Osteoporotic Vertebral Fractures.

Ye C, et al. | *Journal of imaging informatics in medicine* | 2026-10-13

Bone mineral density alone under-predicts osteoporotic vertebral fracture. This source-aware fusion of clinical variables, BMD and lumbar MRI paraspinal muscle representations was tested the hard way: five centers, four external test cohorts, reclassification against the clinical baseline and calibration reported after explicit recalibration.

Key metrics: 2,680 participants over 40 years from five centers (one training, one internal validation, four external test cohorts). Internal validation AUC 0.881 (95% CI 0.854 to 0.909); external 0.851 to 0.917; AUC gain over the clinical plus BMD baseline 0.046 to 0.093 (DeLong $P \leq 0.034$); continuous NRI 0.268 to 1.034. After recalibration, Brier 0.140, calibration slope 0.926; inference 57.71 +/- 18.66 ms. Retrospective.

Clinical relevance: The fracture-prediction gain over BMD alone is consistent rather than dramatic, and the methodological completeness is the story: external multi-site testing, improvement measured against the clinical baseline rather than against nothing, and published calibration numbers. Class activation maps pointing at the paraspinal muscles support the biological plausibility of the signal.

PMID 42870160 DOI 10.1007/s10278-026-02362-8

Automated Lewis Score Estimation for Small Bowel Crohn Disease Using Capsule Endoscopy: Retrospective Cohort Study.

Meng Y, et al. | *JMIR medical informatics* | 2026-10-13

Automated Lewis scoring for small bowel Crohn disease from capsule endoscopy looks feasible: a multitask system combining YOLO-style lesion detection, stenosis classification and LSTM temporal regression tracked expert-consensus scores, with the clear caveat of a small scoring cohort.

Key metrics: 814 patients for lesion-level annotation; 89 cases with expert-consensus Lewis scores (49 training, 40 internal test). Internal test: mean absolute error 2.9 (95% CI 2.1 to 3.8), R2 0.82 (0.70 to 0.90); lesion mean average precision 0.94 at 15 ms per frame. Same-center temporal cohort of 20 consecutive 2026 cases: mean absolute error 3.3; the authors flag these temporal estimates as requiring cautious interpretation. Pediatric center, retrospective.

Clinical relevance: Lewis scoring is slow and reader-dependent, so automation has genuine workflow value, but 40 internal test cases and a 20-case same-center temporal check are a feasibility signal, not deployment evidence. The authors state the requirement themselves: larger multicenter prospective validation with independent external annotation before clinical use.

PMID 42867883 DOI 10.2196/92416

Validation of a Deep Learning Model for Automated Ureteral Length Measurement on CT Against Gross Specimens.

Wang K, et al. | *Journal of imaging informatics in medicine* | 2026-10-13

A deep learning pipeline measured ureteral length on CT and was checked against the hardest reference available, the surgical specimen: the AI measurement was statistically indistinguishable from the gross-specimen reference, while each single urologist measurement differed from it.

Key metrics: Training on 3,629 CT series from 1,502 patients; independent prospective hold-out of 70 radical nephroureterectomy patients with twice-measured gross specimens. AI 269 +/- 2.75 mm vs reference 269 +/- 3.45 mm ($p > 0.999$); each single manual measurement differed from the reference (both $p = 0.002$). Small systematic underestimation in longer ureters (coefficient -0.54 mm per mm of reference length, $p < 0.001$). Single center.

Clinical relevance: The appeal is the reference standard: a physical specimen rather than another reader's opinion, measured in a prospective cohort. Directly useful for preoperative planning and device sizing in urology; the single-center design and 70 prospective cases keep it short of routine deployment, as the authors acknowledge.

PMID 42870153 DOI 10.1007/s10278-026-02400-5

Repurposing adult AI radiograph models for neonatal care through continual learning: an international multicentre study.

Rao VM, et al. | *Nature biomedical engineering* | 2026-10-13

Fewer than one in five commercial imaging AI models are cleared for pediatric use, and adult-trained models generalize poorly to neonates. Before you point an adult-cleared tool at a population it was never validated on, this study shows what a fix looks like: 30 neonatal ICUs in 20 countries each contributed 50 radiographs for sequential continual learning, with no data pooling.

Key metrics: Neonatal chest radiographs from 30 NICUs across 20 countries and territories; 50 radiographs per hospital for sequential retraining of an adult endotracheal tube assessment model. Anatomical landmark localization error 0.67 cm with continual learning vs 1.29 cm for hospital-specific fine-tuning and 4.30 cm for the adult-only model, lower at 29 of 30 sites; improved detection of deep intubations at 27 of 30 sites. Retrospective.

Clinical relevance: An adult-validated model used on neonates is outside its validated population, whatever the installation technically allows. For your AI committee this is a validation and risk-management item under EU AI Act Article 9: inventory which deployed tools see patients outside their cleared or validated population, document that as a risk, and either restrict routing or require an adaptation pathway like this one, with local acceptance data.

PMID 42869977 DOI 10.1038/s41551-026-01812-6

SECTION 2

INDUSTRY & REGULATION

Sources: Official registers | Press releases | FDA databases

[FDA] SimBioSys TumorSight Risk cleared as first MRI-derived breast cancer recurrence prognostic

SimBioSys | 15 October 2026

Cleared for what: 510(k) clearance as a prognostic indicator of recurrence-free survival at five and ten years in adult women with HR-positive, HER2-negative stage I to IIIA breast cancer, computing a 0 to 100 score and a low, intermediate or high risk class from routine breast MRI, pathology and clinical data; labeling says results must be weighed alongside other clinicopathologic factors. Evidence level: vendor materials citing a multicenter study of over 2,100 patients presented at ASCO 2026, where each 10-point score increase was associated with 22% higher recurrence risk; no peer-reviewed publication of the validation is cited in the release. Questions for the vendor: which cohorts trained and which validated the score, and how closely do they match your MRI protocols and population; what are discrimination and calibration by stage and age; what does the score add to, or how does it disagree with, the genomic assays already guiding the same endocrine and chemotherapy decisions; is the threshold locked. Monitoring regardless of clearance: five and ten year endpoints cannot be checked locally for years, so define a surrogate monitoring plan now, starting with score distribution by scanner and site.

Source: [SimBioSys press release](#)**[REGULATION] Thirona LungQ MP CE-certified Class IIb for automated mucus plug analysis on chest CT**

Thirona | 15 October 2026

Certified for what: CE marking under EU MDR as a Class IIb device for fully automated quantification of airway-occluding mucus plugs on chest CT in LungQ version 5, reporting count, volume and location for patient phenotyping and for monitoring mucus burden under treatment; the company claims first-in-class status for this analysis. Evidence level: company announcement; MDR certification attests conformity assessment, not a published reader study, and the release cites a European rollout starting in Spain, the Netherlands and France rather than clinical outcome data. Questions for the vendor: what published validation supports plug detection against expert consensus; how reproducible are the quantitative outputs across CT vendors, reconstruction kernels and slice thicknesses; where is the intended-use boundary between phenotyping research and individual treatment decisions; what change control applies within version 5. Monitoring regardless of certification: quantitative CT biomarkers drift with acquisition changes, so re-baseline on a fixed scan set after every scanner or protocol update.

Source: [Thirona press release](#)**[PARTNERSHIP] Cortechs.ai brings NeuroQuant suite to the CARPL.ai marketplace**

Cortechs.ai / CARPL.ai | 14 October 2026

Announced: a distribution partnership putting Cortechs.ai's quantitative neuroimaging portfolio, NeuroQuant and variants for MS, brain tumor, CT, PET and lesion surveillance plus OnQ Prostate, on CARPL.ai's marketplace, which claims over 300 applications from more than 100 vendors behind one integration, procurement and monitoring channel. What it changes for buyers: platform access lowers integration cost, not the evidence bar; each tool keeps its own regulatory status and validation record. Questions before relying on it: does platform-level monitoring expose the per-tool, per-site metrics your governance needs rather than uptime statistics; who holds deployer obligations under the EU AI Act when a marketplace sits between vendor and hospital; and does the contract let you export the monitoring data when you leave the platform. Local acceptance testing remains yours either way.

Source: [Cortechs.ai press release](#)

[MARKET] Sirona Medical raises \$40M Series D to push agentic AI into the radiology workflow*Sirona Medical | 15 October 2026*

Announced: a \$40 million Series D led by Avidity Partners, total funding around \$150 million, to expand the cloud-native RadOS platform and commercialize Cortex, an agentic AI orchestration layer the company says will coordinate diagnostic models, generate context-aware preliminary interpretations and automate report drafting inside the radiologist's workflow. Evidence level: funding announcement and company description; no clinical performance data accompany the claims. Why it matters to buyers: an orchestration layer that drafts preliminary interpretations is itself a clinical function, not plumbing. Questions to ask: what is the regulatory status of each Cortex capability, especially preliminary interpretation; what evidence supports draft-report quality and error rates; and how does the platform log which model contributed what to a signed report so that accountability survives the convenience.

Source: [Sirona Medical press release \(PR Newswire\)](#)

SECTION 3

POST-MARKET SIGNAL*Drift | Recalls and MAUDE | Real-world monitoring | EU AI Act Article 72***AEGIS: operational infrastructure for post-market governance of adaptive medical AI***13 October 2026*

A npj Digital Medicine paper published October 13 presents AEGIS, an internal governance and post-market monitoring infrastructure designed to make FDA Predetermined Change Control Plans and EU post-market surveillance operational rather than documentary: dataset assimilation and retraining, model monitoring, and a conditional decision module that routes every candidate update into APPROVE, CONDITIONAL APPROVAL, CLINICAL REVIEW or REJECT, with an independent ALARM signal for a released model at risk. On real cross-site data, drift detection reached ROC-AUC 1.00 on batches of about 500 patients with no false alarms, demonstrated on brain tumor segmentation (BraTS) and sepsis prediction. Limits the authors state: proof of concept, covariate shift only, concept drift left for future work. What a deployer can borrow today: the four-outcome taxonomy and the batch-wise drift gate as the skeleton of an Article 72 monitoring file.

Link: <https://www.nature.com/articles/s41746-026-03290-1>**Drift evidence, read critically: what a department can actually watch after go-live***14 October 2026*

A WAIRR explainer published October 14 by Dr Marcus Judge works through the drift literature a UK department would cite, and what each study can and cannot support. The Grampian mammography analysis: a manufacturer threshold set elsewhere would have flagged 48.3% of exams against 5.4% for the first human reader, site calibration brought it to 13.0%, and a scanner software upgrade then roughly tripled the flag rate. A Mayo Clinic CTPA disagreement monitor over 11,999 studies: a 2 to 3% fall in specificity produced a two to threefold rise in discrepancies while a 10% sensitivity drop was needed for the same signal, so disagreement monitors catch extra false alarms quickly and missed findings slowly. The practical conclusion matches the RCR's March 2026 guidance: set a local baseline, watch flag rate and input characteristics weekly, audit samples that include AI-negative cases, and treat every scanner change and model update as a change to the tool.

Link: <https://wairr.co.uk/learn/model-drift-and-monitoring/>

SECTION 4

PLAYBOOK SNIPPET

One concrete step your AI committee can take this month

EU AI Act Article 14 (human oversight)

Step: Before a new tool goes live, record each reader's baseline on a fixed local case set: reading time plus detection or marking rate. Repeat at 3 months with the tool on and report per-reader deltas, never only the pooled average.

Why: Pooled averages hide what a tool does to specific humans: assistance can lift junior readers, slow senior readers, or move two readers in opposite directions, and human oversight is only real if you know which is happening in your own reading room.

Worked example: Zhou Y et al. (Journal of Imaging Informatics in Medicine, PMID 42870161, DOI 10.1007/s10278-026-02352-w) show why the split matters: with an interpretable MRI model for endometrial stromal sarcoma, the junior radiologist's specificity rose from 46.4% to 82.1%, exactly the per-reader delta an Article 14 oversight file should contain. This week's mammography scenario study (PMID 42867728) makes the same point at policy level: the integration scenario, not the model, produced the gain.

SECTION 5

MEDIA HIGHLIGHTS

AuntMinnie | Radiology Business | The Imaging Wire | Diagnostic Imaging | ITN

SCARLET-1: prostate MRI AI matches radiologists on average, varies by site

AuntMinnie.com | 13 October 2026

AuntMinnie covered the SCARLET-1 study published October 13 in Radiology: 22,445 prostate MRI exams from 21,252 men across 46 cities in 22 countries, with external testing on 1,974 patients. The AI reached 83 to 88% agreement with radiologist readings for Gleason grade group 2 or higher cancer in diagnostic and screening settings, with consistent subgroup performance (mean AUC difference 0.02), but per-site AUC spread widely: 28% of sites below 0.80 and 32% at 0.90 or above, with image quality moving AUC between 0.78 and 0.88. The authors' deployment conditions read like a governance checklist: careful operating-point selection, local testing, longitudinal performance monitoring and prospective evaluation of clinical, equity and cost outcomes.

Link: <https://www.auntminnie.com/clinical-news/mri/news/15837355/ai-on-par-with-radiol...>

Neiman HPI pushes back on claims that AI has already reduced the need for radiologists

AuntMinnie.com | 13 October 2026

A JACR editorial by Harvey L. Neiman Health Policy Institute leaders disputes recent JAMA Health Forum comments that falling imaging use let the US meet demand without more radiologists. The HPI team cites its own projections of 17 to 27% growth in imaging use by 2055, Medicare data showing CT and MRI volume growth since 2016, radiologist workload in RVUs per beneficiary at its highest recorded level, and a 48.3% inflation-adjusted cut in radiologist reimbursement between 2005 and 2025. Their concern: optimistic AI-workforce claims from outside the field could repeat the chilling effect of the 2016 predictions and undermine workforce planning while workload keeps rising.

Link: <https://www.auntminnie.com/imaging-informatics/artificial-intelligence/article/1...>

A three-step buyer's guide for radiology AI, with a risk matrix for downstream effects

Radiology Business | 15 October 2026

Radiology Business profiled a JIIM open-access paper by Baird and colleagues, including ACR board chair Christoph Wald, proposing a three-step evaluation for any radiology AI purchase: establish how the model will perform on the local population, map how often its errors would change local clinical decisions using a structured matrix the authors call RADDE, then weigh error frequency against error cost. The authors note that no independent comparative performance analysis of imaging AI detection models, in the style of Consumer Reports, has ever been published, which leaves vendor narratives unchallenged unless practices run this kind of structured local evaluation themselves.

Link: <https://radiologybusiness.com/topics/artificial-intelligence/3-step-buyers-guide...>

SECTION 6

PROFESSIONAL SOCIETIES

SIIM | ACR | RCR | Official publications and event coverage

SIIM webinar demonstrates autonomous chest X-ray AI and its validation thresholds

14 October 2026

SIIM's Enterprise Imaging webinar series hosted Sectra on October 14 for a live demonstration of autonomous AI clearing normal chest exams from the radiologist worklist, the capability acquired with Oxipit, alongside third-party pathology AI on the Amplifier marketplace. The session's stated focus included the validation and safety thresholds applied before deploying autonomous AI; the algorithms are CE-marked, with FDA submission in process. Autonomous normal-clearing is the sharpest test yet of local validation discipline, since no human reads what the model clears.

Source: <https://siim.org/webinar/ai-in-practice-autonomous-radiology-ai-and-ai-supported-pathology/>**RSNA reports its Imaging AI Certificate Program is building evaluation and monitoring skills**

16 October 2026

RSNA published participant feedback on its Imaging AI Certificate Program on October 16: radiologists and imaging professionals who completed the Foundational, Advanced, Emergency or Chest certificates report stronger skills in evaluating, implementing and monitoring AI tools, with a Neuroradiology certificate scheduled for late 2026. Structured, role-specific AI training is one of the few concrete responses available to the EU AI Act's Article 4 literacy obligation and to equivalent competence expectations outside Europe.

Source: <https://www.rsna.org/news/2026/october/rsna-ai-certificate-program-success>**JIIM publishes SIIM 2026 abstract supplement on deploying AI under governance constraints**

14 October 2026

The Journal of Imaging Informatics in Medicine released its Volume 39 supplement with the scientific abstracts of the SIIM 2026 Annual Meeting, announced in SIIM's October 14 NewsBrief. The society highlights work on AI deployment with governance, privacy and interoperability constraints, plus workflow optimization and multimodal integration, a useful scan of what imaging informatics teams are actually building this year.

Source: <https://multibriefs.com/briefs/siim/SIIM101426.php>

SECTION 7

KEY OPINION LEADERS

LinkedIn | Public statements | Week of 11 to 17 October 2026

Pranav Rajpurkar

Co-founder of a2z Radiology AI; Associate Professor, Harvard

Announced the Nature Biomedical Engineering paper behind this week's flagship selection: 75 authors, 30 neonatal ICUs in 20 countries, each contributing just 50 chest X-rays through the MAIDA initiative to adapt an adult AI model for checking breathing-tube placement in newborns, with accuracy improving at 29 of the 30 hospitals. His framing: newborns are among the patients medical AI serves least, and small, federated local contributions can fix that without data sharing.

Source: LinkedIn, 15 October 2026

Woojin Kim

CMIO, ACR Data Science Institute; Chief Strategy Officer and CMIO, HOPPR

Shared the ACR Data Science Institute's wrap-up of EuSoMII 2026 in Crete, where he presented the ACR's initiatives on imaging AI quality, implementation and performance monitoring, including the Assess-AI registry, alongside ACR Informatics Committee colleagues Alex Towbin and Ali Tejani. The through-line of the institute's message: the work on imaging AI quality continues well after deployment.

Source: LinkedIn, 13 October 2026

Amine Korchi*Radiologist; HealthTech Innovation and Ventures*

Argued that many radiology inefficiencies share one root cause: poor cross-functional understanding between radiologists, technologists and secretaries, illustrated by a follow-up MS MRI where prior images sat in PACS unlinked to the patient and invisible to the reader. His conclusion: operational excellence requires institutional introspection, well-designed processes, standard operating procedures and continuous staff training, not more individual effort. A reminder that AI lands on top of exactly these processes.

Source: [LinkedIn](#), 15 October 2026

QUALITY ASSURANCE NOTE

Compiled by Dr. Sergey Morozov from publicly available sources: peer-reviewed papers indexed in PubMed (PMIDs and DOIs verified, no preprints or arXiv), official press releases, regulatory databases and specialist media. It does not constitute medical, regulatory or financial advice. All papers have publication dates verified within 11 to 17 October 2026. Paper selection is deterministic (pinned PubMed query, Q1/flagship journal allow-list, exclusion of pure reviews [meta-analyses kept], radiomics-only and interventional-radiology work, then a transparent ranking score). Industry items come from official press releases / regulatory notices; KOL items from public LinkedIn activity within the window.