

WEEKLY AI INTELLIGENCE REPORT

Radiology and Medical Imaging

Week 40 | 27 September to 3 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

5

Industry / FDA items

3

Media highlights

3

Key opinion leaders

GOVERNANCE TAKEAWAY OF THE WEEK

Stop evaluating AI tools on average performance and start evaluating them on where the average breaks: this week supplies the instruments to do it. Put two new lines in your pre-deployment template, performance by subgroup (with missing subgroup evidence logged as a risk, not a footnote) and end-to-end latency against the report-finalization window per modality. The JMIR refined-exclusion gates and the NHS 8-domain checklist give you the structure; the 20-center latency study and the breast AI agreement data show exactly what those lines will catch.

SECTION 0

EXECUTIVE SUMMARY

Key trends, week of 27 September to 3 October 2026

[GOVERNANCE] THE TOOLING FOR LOCAL AI GOVERNANCE IS MATURING

Two publications hand committees concrete instruments this week: a JMIR viewpoint defines 'refined exclusion' and proposes lifecycle gates with a four-action vocabulary (proceed, enrich, restrict, retire), and an NHS trust publishes an 8-domain preimplementation checklist derived from its own fracture-AI and prostate-MRI deployments. Both convert governance from prose into auditable steps.

[EVIDENCE] WHAT CREDIBLE ULTRASOUND AI VALIDATION LOOKS LIKE

Zhu Y et al. (Radiology: AI) validate a dual-modality ultrasound video model for superficial lymphadenopathy across five centers with a prospective external video set: AUC 0.91 to 0.95, subgroup AUCs 0.85 to 0.93, and junior-reader AUC lifted from 0.72 to 0.88 with assistance. Registered, external, prospective: the bar vendors should be held to.

[IMPLEMENTATION] INFRASTRUCTURE, NOT ALGORITHMS, IS THE BOTTLENECK

A 4.5-year, 20-center, 96,874-exam observational study finds 72% of AI latency comes from data routing, and the AI result arrives after report finalization in 7.2% of cases (13.2% for chest CT). Turnaround gains concentrate in high-volume modalities; sentiment declined where latency bit. Measure end-to-end latency per modality before and after any deployment.

[OVERSIGHT] CO-REPORTED AI SCORES ARE NOT ONE TEST

In a biopsy-referred breast cohort, a platform's tomosynthesis and synthesized-2D malignancy scores showed similar AUC but only moderate agreement (kappa 0.53), flagging different cancers. Validate and threshold each output separately, and never downgrade a suspicious finding on a low AI score alone.

[WATCHPOINT] AVERAGE PERFORMANCE CONCEALS LOCAL FAILURE

The week's thread runs from theory to plumbing: a governance framework for safety concentrated in unmeasured subgroups, a checklist built from real deployment gaps, a latency study showing where AI results miss the report window, and a breast cohort where two scores from one platform disagree on which cancers to flag. Aggregate accuracy is the wrong unit of decision; your population, your scanners and your pipes are the right one.

SECTION 1

PEER-REVIEWED PAPERS

Source: PubMed | Verified indexing | No preprints | 27 September to 3 October 2026

Reducing Clinician Annotation Fatigue in Open-Set Federated Learning: A Client-Adaptive Vision-Language Gatekeeper.

Nesturi A, et al. | *Journal of the American College of Radiology* | 2026-10-01

Before committing radiologist hours to annotation rounds in a federated learning project, put a quality gate on the annotation queue itself. This JACR study adds a client-adaptive vision-language gatekeeper (PromptGate) that withholds artifact-containing mammograms unsuitable for a breast-density labelling task, tested on FedEMBED (222,700 training and 65,891 test images across four scanners, artifact prevalence 4.8% to 13.1%) with a blinded, randomized 500-view reader cohort.

Key metrics: Withheld 101 of 154 artifact images and 54 of 346 relevant images: artifact sensitivity 65.6%, in-distribution retention 84.4%, false-removal rate 15.6%. Gross annotation time fell about 30% in the enriched cohort (3.82 to 2.68 hours); crediting only correctly removed artifacts, 20.0%. Artifact discrimination AUC 0.85. Queue purity 92.5% versus 88.5% unfiltered and 87.0% with a static filter; reader agreement with the manual artifact reference 93.6% (kappa 0.85), with the clinical-record density label only 60.9% exact (kappa 0.48). Retrospective, single-reader, proof of concept.

Clinical relevance: The striking secondary finding is the weak agreement between an experienced reader and the clinical-record density labels (kappa 0.48): the reference labels your federated model learns from may be the weakest link. For your AI committee this is a data-governance item under EU AI Act Article 10: any project that consumes local annotations should document queue-curation rules and label reliability the same way it documents cohort composition, and acceptance testing of a training pipeline should include an audit of label quality, not only model output.

PMID 42822685 DOI 10.1016/j.jacr.2026.09.023

Deep Learning Analysis of Dual-Modality US Videos for the Characterization of Superficial Lymphadenopathy.

Zhu Y, et al. | *Radiology. Artificial Intelligence* | 2026-09-30

If you want a template for how ultrasound video AI should be validated, read this one. DMUVL-DiagNet integrates B-mode and color Doppler video with basic clinical data to separate benign from malignant superficial lymphadenopathy and to split lymphoma from metastasis, trained on 1,616 retrospective videos and tested externally on a prospective 454-video set and 1,956 static images across five centers, with a registered protocol (ChiCTR2400090592) and an AI-assisted reader study.

Key metrics: 4,026 patients overall (median age 52.6 years). AUC 0.91 to 0.95 for benign versus malignant and 0.87 to 0.91 for lymphoma versus metastasis across all test sets; subgroup AUCs 0.85 to 0.93. The model exceeded the average of every radiologist experience group. With AI assistance, junior radiologist AUCs rose from 0.72 to 0.88 (benign versus malignant) and from 0.58 to 0.77 (lymphoma versus metastasis), all $P < .001$.

Clinical relevance: Kept on clinical merit. Superficial lymphadenopathy triage is a daily, operator-dependent task, and the largest gain here is not the model's standalone AUC but the lift it gives junior readers on the harder lymphoma-versus-metastasis call, which changes who gets biopsy and who gets flow cytometry first. The prospective external video set and the reported subgroup range are what make the numbers credible.

PMID 42813999 DOI 10.1148/ryai.260131

AI Latency, Report Turnaround Time, and Adoption in a Multi-Vendor AI Ecosystem: A Multi-Site Observational Study.

Morozov S, et al. | *Journal of the American College of Radiology* | 2026-09-29

Before you benchmark another algorithm, benchmark your pipes. This 4.5-year observational study of 10 AI tools from 7 vendors across a 20-center outpatient network measured where AI results actually arrive too late to matter and what that does to turnaround time and radiologist sentiment, over 96,874 examinations.

Key metrics: Median PACS-to-PACS latency 2.06 minutes (IQR 1.74 to 3.05), 72% of it data routing rather than inference. The AI result arrived after report finalization in 7.2% of cases overall, 3.0% for knee MRI and 13.2% for chest CT. In 20,909 examinations, AI availability was associated with lower median turnaround time for trauma radiography (-26%) and knee MRI (-18%, both $p < 0.001$); brain volumetry showed no change (+9.2%, $p = 0.33$). Adoption 91.4% of 58 radiologists; NPS declined for chest CT (+38 to -3) and aorta CT (+22 to -25), nominally significant before correction. Program cost about 23% to 25% of one radiologist FTE, generating 0.69 FTE of capacity through trauma radiography alone (0.46 after sensitivity analysis).

Clinical relevance: The limiting factor at scale was infrastructure latency, not algorithm quality, and the 'too late' rate varied fourfold by modality. For your AI committee this changes both acceptance testing and monitoring: acceptance tests should measure end-to-end latency against the report-finalization window per modality, and the too-late rate belongs on the monthly monitoring dashboard alongside concordance, mapping directly to EU AI Act Article 72 post-market monitoring of whether the system delivers its intended benefit in your workflow.

PMID 42810607 DOI 10.1016/j.jacr.2026.09.026

Tomosynthesis and Synthesized-2D Breast AI Outputs in a Biopsy-Referred Cohort: A Diagnostic Study of Agreement and Incremental Decision-Support Value.

Şahin E, et al. | *Journal of imaging informatics in medicine* | 2026-09-29

If your breast AI platform co-reports a tomosynthesis score and a synthesized-2D score, treat them as two different tests, not one. In 139 biopsy-referred breast-level observations with histopathology reference, the two malignancy scores from the same deployed platform showed similar overall discrimination but disagreed substantially within individual breasts.

Key metrics: AUC 0.728 (tomosynthesis output) versus 0.717 (synthesized output), $p = 0.786$, yet agreement was only moderate (kappa 0.53): among malignant breasts, 21 were flagged by the tomosynthesis output alone versus 6 by the synthesized output alone ($p = 0.006$). Adding either score to a clinical-radiological baseline model improved fit ($p < 0.001$) but moved AUC only from 0.809 to 0.841-0.848. Of 30 cancers negative by the tomosynthesis output, 25 had no mammographic correlate on masked re-review. Single center, retrospective, biopsy-referred cohort.

Clinical relevance: The operational rule the authors derive is the one to adopt now: never use a low AI score by itself to downgrade a suspicious finding, and when a platform emits multiple outputs, validate and threshold each separately. For your AI committee this is an acceptance-testing item under EU AI Act Article 15 and a human-oversight instruction under Article 14: the local acceptance protocol should state which of the co-reported outputs is in clinical use, at what threshold, and what readers are told about the other.

PMID 42811225 DOI 10.1007/s10278-026-02255-w

Refined Exclusion in Medical AI: Data Justice and Patient Safety Governance.

Lee JH, et al. | *Journal of medical Internet research* | 2026-10-01

Before your next procurement or renewal decision, ask not whether the model performs well on average but for whom its safety is undemonstrated. This JMIR viewpoint names the failure mode 'refined exclusion': a system that looks successful in aggregate while uncertainty, error or reduced reliability concentrates in populations that are insufficiently represented, measured, validated or monitored.

Key metrics: Viewpoint article: no quantitative results, by design. It distinguishes model-level disparity, patient safety signals and documented harm, drawing on illustrative evidence from chest radiography, dermatology, computational pathology and medical foundation models, and proposes lifecycle decision gates covering intended use, subgroup learnability, data provenance, validation, procurement, deployment, monitoring, updates and patient feedback, each linking minimum evidence to one of four actions: proceed, enrich or validate, restrict use, or pause or retire.

Clinical relevance: The usable part is the gate structure: every gate names who decides and what the minimum evidence is, which is precisely what most hospital AI policies lack. For your AI committee this changes validation and monitoring at once: adopt the four-action vocabulary (proceed, enrich, restrict, retire) in your review template, and treat missing subgroup evidence as a tracked risk under EU AI Act Article 9 and a data-governance gap under Article 10 rather than a footnote.

PMID 42821994 DOI 10.2196/102359

Dual-discriminator Generative Adversarial Network (DD-GAN) for Virtual Elastography Ultrasound in Thyroid Cancer Diagnosis.

Lu W, et al. | *Academic radiology* | 2026-09-28

Elastography without an elastography unit: a dual-discriminator GAN generates virtual elastography images directly from conventional B-mode thyroid ultrasound, aimed at portable and operator-limited settings where real strain imaging is unavailable or unreproducible.

Key metrics: Diagnostic performance of the virtual strain ratio was statistically indistinguishable from real elastography: AUC 0.744 versus 0.774 internally ($p = 0.318$), 0.742 versus 0.759 on external testing ($p = 0.461$) and 0.751 versus 0.761 in prospective validation ($p = 0.772$). Radiologists could not tell virtual from real images in blinded testing, and adding the virtual-elastography Tsukuba score to TI-RADS significantly improved AUC for both junior and senior readers.

Clinical relevance: Kept on clinical merit. The honest reading is that elastography itself is a modest discriminator (AUCs around 0.75) and the synthetic version faithfully reproduces that modesty, including in a prospective set. The value is access: TI-RADS plus a stiffness signal on hardware that has no elastography mode, which matters for point-of-care and resource-limited thyroid workup more than for well-equipped departments.

PMID 42805865 DOI 10.1016/j.acra.2026.09.019

Evaluating the Long-Term Malignancy Risk Prediction Accuracy of Major Guidelines for Pancreatic Cystic Lesions Using Radiology Features.

Liu T, et al. | *Journal of the American College of Radiology* | 2026-09-28

A rare long-horizon test of the guidelines themselves: using a previously validated large language model to extract pancreatic cystic lesion features from 25 years of radiology reports, the authors scored 13,269 patients against ACR, Kyoto and European guidelines and followed them for 5-year malignancy.

Key metrics: 209 of 13,269 patients (1.6%) developed malignancy within 5 years. All three guidelines showed modest discrimination: AUROC 0.655 (Kyoto), 0.639 (ACR), 0.623 (European), $p < 0.001$ for the comparison. High-risk classification was fully concordant across guidelines, but worrisome-feature assignment varied from 308 to 507 patients at similar 5-year cancer incidence (10.3% to 11.7%). After 5 years of imaging stability ($n = 1,830$), 8-year cancer incidence was 0.71% overall and 0.46% in patients meeting Kyoto discontinuation criteria. Single center, retrospective.

Clinical relevance: Kept on clinical merit. For surveillance practice the actionable number is the 0.46% 8-year incidence after five stable years under Kyoto discontinuation criteria, which supports ending surveillance in that group, and the modest AUROCs are a reminder that imaging features alone carry limited long-term predictive power however the guideline weights them.

PMID 42805547 DOI 10.1016/j.jacr.2026.09.019

Translating Real-World Safety and Implementation Gaps Into a Deployment-Derived AI Readiness Preimplementation Checklist for NHS Health Care Providers: Checklist Development Study.

Adesuyi A, et al. | *JMIR AI* | 2026-09-28

Before your next go-live, compare your pre-implementation paperwork with this one: an NHS district general hospital distilled its real deployment experience, an AI fracture detection system and an AI-supported prostate MRI pathway, into the SID & ADE preimplementation readiness checklist, then applied it retrospectively to the same deployments to see what it would have caught.

Key metrics: Checklist of 8 domains: use-case definition; clinical safety and accountability; local validation and performance; workforce readiness and human factors; operational and technical integration; information governance and ethics; procurement, liability and financial risk; monitoring, evaluation and stop rules. Evidence sources included a clinico-AI discordance study and plan-do-study-act cycles. Retrospective application showed recurrent gaps in workforce readiness, local validation and monitoring. No quantitative performance data; face and content validity claimed only for the originating trust.

Clinical relevance: Deployment-derived beats committee-derived: the recurrent gaps it surfaces (workforce readiness, local validation, monitoring and stop rules) are exactly the domains generic frameworks leave vague. For your AI committee this is an acceptance-testing instrument: map its 8 domains onto your current pre-go-live sign-off and record which domains your process does not yet cover, noting that stop rules and monitoring correspond to EU AI Act Articles 14 and 72 obligations you will have to evidence anyway. Independent prospective validation is still pending, so treat it as a deliberation aid, not a certified standard.

PMID 42804536 DOI 10.2196/93900

SECTION 2

INDUSTRY & REGULATION

Sources: Official registers | Press releases | FDA databases

[FDA] Visiana's BoneXpert bone-age software cleared for the US market

Visiana | 30 September 2026

Cleared for what: 510(k) clearance (K262390) to support pediatric radiologists in assessing bone age from hand radiographs in patients aged 2 to 21 years by the Greulich-Pyle method; it is an aid, not an autonomous read. On what evidence: the company cites multi-site accuracy documentation from pediatric radiologists at five US sites and longitudinal consistency work; BoneXpert has carried a CE mark and been in European clinical use since 2009, so a long real-world record exists outside the US. Questions for the vendor: per-site accuracy at the five US sites and the demographic composition of that cohort; performance in the age extremes (2 to 3 and 18 to 21 years) and across skeletal dysplasias; test-retest consistency figures, since bone-age change between visits is often the clinical question. Regardless of clearance, set a monthly concordance sample against your own readers before letting the number flow into endocrinology workflows.

Source: [Visiana press release](#)**[FDA] DeepHealth clears a foundation-model chest X-ray CADe**

DeepHealth | 29 September 2026

Cleared for what: 510(k) clearance for Chest XRay, a computer-aided detection (CADe) tool that detects and localizes suspected thoracic abnormalities on chest radiographs; one of the first US clearances for a chest X-ray CADe built on a foundation model. On what evidence: company release; the underlying model runs in a CE-marked version the vendor says analyzes over 2.8 million examinations a year in Europe (internal deployment data, not a published study). Questions for the vendor: the pivotal reader-study data behind the US clearance, finding-level sensitivity and false-positive rates per image, which specific abnormality classes are inside the cleared claim, and what happens to the claim when the foundation model is updated: is there a Predetermined Change Control Plan, and what does it commit to in postmarket monitoring. A broad 'thoracic abnormalities' label makes local acceptance testing on your own casemix more important, not less.

Source: [Axis Imaging News](#)**[FDA] CapsoVision's AI Highlights cleared on a competitor's predicate**

CapsoVision | 28 September 2026

Cleared for what: 510(k) clearance (letter dated 25 September, announced 28 September) for AI Highlights, which marks and prioritizes frames with suspected abnormal lesions in small-bowel capsule endoscopy performed for suspected small-bowel bleeding, in adults and children over two; it is not validated for other indications such as Crohn's surveillance. On what evidence: an unpublished retrospective reader study; no sensitivity, specificity or reading-time figures have been published. Notable: the predicate is a rival's device, AnX Robotica's NaviCam ProScan, which itself needed De Novo authorization in 2023 because no predicate existed; CapsoVision rode the classification that decision created. Questions for the vendor: the reader-study design and endpoints, per-lesion-type detection rates, and the miss rate when physicians review only highlighted frames. Until performance data are published, treat reading-time claims as unverified.

Source: [CapsoVision press release \(GlobeNewswire\)](#)**[MARKET] Stockholm Region adopts AI as the independent second reader in breast screening**

Lunit | 28 September 2026

What was announced: Lunit INSIGHT MMG will replace the radiologist second read across the entire Stockholm Region (Capio Saint Göran, Karolinska University Hospital, Södersjukhuset), deployed with Sectra, expected to cover 200,000 to 250,000 examinations a year; it is Lunit's second region-level deployment in Sweden after Dalarna. On what evidence: unusually strong for a deployment announcement; ScreenTrustCAD, a prospective population-based study of over 55,000 women at Saint Göran, plus roughly 200,000 routine examinations since 2023, through changes of mammography systems. Questions for any site considering the same move: what the monitoring plan tracks once the human second reader is gone (interval cancers, recall rate by subgroup, drift after equipment changes), who holds authority to suspend the AI reader, and what the fallback staffing plan is. Replacing a reader is a different risk class from assisting one; the monitoring obligation rises accordingly.

Source: [Lunit press release](#)

[MARKET] Clairity Breast goes direct-to-consumer via Everlywell*Clairity* | 1 October 2026

What was announced: nationwide US consumer access to Clairity Breast, the AI tool that estimates five-year breast cancer risk from a routine screening mammogram, through Everlywell's platform with ordering and result release by Everlywell-affiliated clinicians. Regulatory status: FDA De Novo authorization (May 2025), meaning no predicate device existed; it predicts future risk and does not detect cancer. Evidence: validated in more than 77,000 screening mammograms at five US sites for authorization and over 120,000 across ten facilities after; now cited as an option in NCCN screening guidance for risk assessment from age 35. Questions radiology leaders should ask: how a consumer-ordered risk score re-enters the clinical record, who counsels the patient when the score conflicts with the local risk assessment, and whether downstream supplemental-screening demand (MRI, ultrasound) has been modelled. A DTC channel moves the governance burden from the imaging center to whoever receives the downstream referrals.

Source: [Everlywell press release](#) (Business Wire)

SECTION 3

POST-MARKET SIGNAL*Drift* | *Recalls and MAUDE* | *Real-world monitoring* | *EU AI Act Article 72***FDA's genAI discussion paper would shift the evidence burden to post-market: comments close 19 October***28 September 2026*

Legal analyses published this week (Cooley, 28 September; Haynes Boone, 2 October) unpack the FDA CDRH discussion paper on generative-AI-enabled devices and converge on the same point: CDRH is openly considering accepting greater premarket uncertainty in exchange for heavier reliance on post-market monitoring, via periodic benchmarking against premarket thresholds, sample-based independent clinician review of real-world inputs and outputs, and performance-degradation monitoring, and even asks whether machine-based supervisory agents could do part of that oversight. Commenters so far are pushing back that post-market monitoring is not an adequate substitute for premarket evidence where harm is severe or irreversible, and want sponsors to state what failure modes a monitoring program cannot detect. Deployers should care because every one of these mechanisms presumes the hospital can supply or host the monitoring data. Feedback closes 19 October 2026 (docket FDA-2026-N-7874); imaging departments with deployed AI have standing to comment and rarely do.

Link: <https://www.cooley.com/news/insight/2026/2026-09-28-regulating-ai-like-a-doctor-...>

SECTION 4

PLAYBOOK SNIPPET*One concrete step your AI committee can take this month*

EU AI Act Article 72 (post-market monitoring); ISO/IEC 42001 performance evaluation clause. The AI Act application calendar is under revision following the Digital Omnibus proposal, so verify dates against the current text

Step: Sample 20 AI outputs per month against the final signed report. Track one number: concordance rate. Plot it. Add a second line this month: the share of AI results that arrived after report finalization, per modality. A trend, not a threshold, is what tells you something has changed.

Why: Drift is invisible in aggregate accuracy and visible in a monthly line. Starting the series now means you have history when the obligation bites; starting later means you have a snapshot.

Worked example: Morozov S et al. (Journal of the American College of Radiology, PMID 42810607, DOI 10.1016/j.jacr.2026.09.026) show what the second line catches: across 96,874 examinations, 7.2% of AI results arrived after the report was signed, ranging from 3.0% for knee MRI to 13.2% for chest CT. A tool whose output misses the reporting window is drifting operationally even when its accuracy has not moved.

SECTION 5

MEDIA HIGHLIGHTS

*AuntMinnie | Radiology Business | The Imaging Wire | Diagnostic Imaging | ITN***Over 1,000 FDA-cleared radiology AI devices, and what the database actually shows***Radiology Business | 29 September 2026*

Radiology Business covered a new Academic Radiology analysis characterizing 1,094 radiology AI devices authorized by the FDA through December 2025. Annual authorizations rose 42% between 2023 and 2025, while the functional makeup of the devices stayed relatively stable. For buyers the useful point is the baseline: the authors position the review as a reference against which the arriving wave of foundation-model and generative tools can be compared, and as a starting inventory check for practices that do not know which of their tools sit in which regulatory category.

Link: <https://radiologybusiness.com/topics/artificial-intelligence/radiology-ai-author...>**Radiology Partners starts docking work units for AI-drafted reports***Ben White (benwhite.com) | 30 September 2026*

Radiologist Ben White dissected Radiology Partners' announcement that from 1 October it will 'share the efficiency gains' of its Mosaic AI drafting platform by reducing time-based work units for exams where AI drafts are available, effectively lowering internal credit per read as read times fall. His analysis is the clearest public account yet of the economic logic AI introduces into radiology employment: vertically integrated groups capture efficiency gains, payers keep paying the same, and the radiologist's work unit absorbs the difference. Whatever one's position, this is post-deployment impact of AI on working conditions, and it will shape how radiologists respond to every future efficiency claim.

Link: <https://www.benwhite.com/radiology/number-goes-down/>**Multi-vendor AI implementation study draws trade coverage***Radiology Business | 1 October 2026*

Radiology Business reported on this week's JACR multi-site observational study of 10 AI tools from 7 vendors across a 20-center outpatient network (also covered in section 1 of this issue), highlighting the finding that infrastructure latency, not algorithm speed, was the primary barrier to clinical utility, and that turnaround-time gains concentrated in high-volume modalities such as trauma radiography. The pickup matters because implementation-economics evidence of this kind rarely reaches the trade press, which usually covers accuracy claims.

Link: <https://radiologybusiness.com/topics/artificial-intelligence/widespread-ai-adopt...>

SECTION 6

PROFESSIONAL SOCIETIES

*SIIM | ACR | RCR | Official publications and event coverage***ACR-SIIM Practice Parameter for Imaging AI takes effect***1 October 2026*

The first ACR-SIIM Practice Parameter for Imaging Artificial Intelligence, adopted by the ACR Council in May 2026, entered into effect on 1 October 2026 under the ACR's standard cycle, by which parameters approved in a given year become effective on 1 October of that year. From this week it is the reference US practice standard for imaging AI: it requires an AI governance group, an inventory of deployed tools with versions and intended use, local acceptance testing before deployment, real-world performance monitoring with defined stop rules, and HIPAA-grade access controls and logging, and it explicitly covers non-FDA-regulated tools including generative AI. Practices aligning with it can pursue the ARCH-AI designation.

Source: <https://www.acr.org/News-and-Publications/Media-Center/2026/first-practice-parameter-for-imaging-ai>

RSNA ships a major update to its AI certificate program

27 September 2026

RSNA's Foundational AI Certificate program is now available in version 2.0, a rebuilt curriculum covering evaluation, implementation and use of imaging AI, alongside the society's modules on emergency triage AI, chest AI and an advanced course on model development and data curation. With EU AI Act Article 4 making AI literacy an obligation rather than an aspiration, a structured certificate is one auditable way to evidence workforce training, though a certificate alone does not demonstrate the tool-specific competence the obligation targets.

Source: <https://radiologybusiness.com/topics/management/education-training/rsnas-ai-certificate-program-undergoes-major-update>

SECTION 7

KEY OPINION LEADERS

LinkedIn | Public statements | Week of 27 September to 3 October 2026

Nina Kottler

Chief Medical AI Officer, Mosaic Clinical Technologies

Posted 2 October a two-month field report answering the question 'what does a Chief Medical AI Officer actually do': travel across US sites, interviews, podcasts, conference presentations and resident teaching, discussions with healthcare leaders on bringing AI into clinical practice, and writing and research on how AI changes clinical decision-making, how to implement it responsibly, and how to monitor it once it is in use, alongside work on AI regulation and reimbursement. A useful concrete sketch of the emerging CMAIO role for any organization deciding whether to create one.

Source: [LinkedIn](#)

Woojin Kim

Chief Strategy Officer and CMIO, HOPPR; CMO, ACR DSI

Posted 2 October an experiment turning his 2024 Radiology editorial 'Seeing the Unseen' into a 25-second animated explainer generated with Anthropic's Claude Opus 5.5, sharing the exact prompt he used. His caveat is the substance: the model built the storyboard and code, but it confidently mislabelled an absolute-washout value and 'seeing what's wrong still takes a human', a compact demonstration of both the speed of generative tools for radiology education and why expert review stays in the loop.

Source: [LinkedIn](#)

Amine Korchi

Radiologist; HealthTech Innovation & Ventures

Posted 2 October ahead of a remote Lunch & Learn at Hospital Israelita Albert Einstein in São Paulo, discussing his recent paper with Jan Beger and Christoph Agten on the future of radiology and radiologists in the AI era, built around three scenarios for the specialty: expansion, concentration and stratification. The international teaching circuit around that framework shows how actively radiology leadership audiences are now working through AI scenario planning.

Source: [LinkedIn](#)

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 27 September to 3 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.