

WEEKLY AI INTELLIGENCE REPORT

Radiology and Medical Imaging

Week 41 | 4 to 10 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

Name the owner. Article 4 of the EU AI Act makes AI literacy an organizational obligation, and this week supplies both the map and the motive: EFOMP's new framework turns AI competence into an auditable professional ladder with six named gaps, and a carotid ultrasound reader study shows why it matters, with AI assistance lifting junior readers from near-chance baselines to usable but still below unassisted seniors. This week, assign one named, competent owner to each deployed model and have them produce a single signed page: what the tool is cleared for, what your acceptance test showed, which patient subgroups and reader groups were checked, and what triggers suspension.

SECTION 0

EXECUTIVE SUMMARY

Key trends, week of 4 to 10 October 2026

[GOVERNANCE] AI COMPETENCE BECOMES A NAMED, AUDITABLE RESPONSIBILITY

EFOMP's updated framework makes AI a core competency across all three medical physics curricula and defines the profession's responsibilities under the EU AI Act in four advanced domains: validation and QA, lifecycle management, generative AI and foundation models, and governance. It is a template any radiology AI committee can borrow to document who is competent to accept and monitor a model.

[EVIDENCE] ASSISTANCE HELPS MOST WHERE OVERSIGHT IS WEAKEST

In a two-center carotid ultrasound stroke-risk study (external AUC 0.939), AI assistance lifted junior readers from near-chance baselines (0.450 and 0.548) to 0.700 and 0.659, while senior readers rose from 0.760 and 0.868 to 0.857 and 0.929. The tool narrows the seniority gap but does not close it: assisted juniors still trail unassisted seniors.

[IMPLEMENTATION] THE DEPLOYMENT PLUMBING IS MATURING

A cloud-native multicenter research platform (edge DICOM nodes, zero-footprint viewer, no VPNs) and an inline-integrable restoration model that brings whole-brain vessel wall MRI under 6 minutes target the same bottleneck: getting validated AI into routine workflow rather than into another retrospective study.

[WATCHPOINT] OPEN FOUNDATION MODELS SHIFT THE VALIDATION BURDEN TO YOU

MedGemma (Nature Medicine) releases open medical vision-language weights with out-of-distribution benchmark gains of 2.6 to 18.1% over base models; UPER unifies diffusion MRI across 17 datasets. Open weights mean anyone can fine-tune, so published numbers are a floor for development, not evidence for your deployment: local validation becomes the deployer's obligation.

[OVERSIGHT] COMPETENCE IS THE SCARCE RESOURCE, NOT MODELS

The week's thread: a professional framework that assigns AI responsibilities to named competencies, a reader study showing assistance cannot substitute for reader skill, and open foundation models anyone can adapt. Models are becoming abundant; the local capacity to accept, monitor and suspend them is not. Fund the people and the audit trail before the next procurement.

SECTION 1

PEER-REVIEWED PAPERS

Source: PubMed | Verified indexing | No preprints | 4 to 10 October 2026

Multimodal Carotid Ultrasound-based Artificial Intelligence Model for Predicting Acute Ischemic Stroke Risk in Patients with Type 2 Diabetes Mellitus.

Sun Y, et al. | *Academic radiology* | 2026-10-09

Before you let an assistive model loose across a mixed reading room, measure what it does to each seniority band separately. This two-center study built a dual-scale Swin Transformer model on carotid ultrasound, fused deep learning features with radiomic and clinical data to predict acute ischemic stroke risk in type 2 diabetes, and then ran a reader study stratified by experience.

Key metrics: 480 patients from two centers (394 development, 86 independent external test). Multimodal AUC 0.952 (95% CI 0.910 to 0.984) internal and 0.939 (0.838 to 0.996) external. With AI assistance, senior reader AUCs rose from 0.760 to 0.857 and from 0.868 to 0.929; junior readers from 0.450 to 0.700 and from 0.548 to 0.659. Diagnostic correctness improved significantly in both groups (GEE $P=0.00235$ senior, $P=0.000343$ junior). Retrospective.

Clinical relevance: The unassisted junior baselines (0.450 and 0.548, one at chance level) are the finding to act on: the model lifts junior readers most, yet their assisted performance still trails unassisted seniors, so assistance is not a substitute for competence. For your AI committee this is an acceptance-testing and oversight item under EU AI Act Article 14: record per-seniority reader baselines and deltas before an assistive tool goes live, because a pooled average will hide exactly this split.

PMID 42855384 DOI 10.1016/j.acra.2026.09.036

Peritumoral ultrasound-driven multimodal AI model for predicting lymphovascular space invasion and prognosis in cervical cancer.

Shen H, et al. | *Abdominal radiology (New York)* | 2026-10-08

A prospective test cohort, not a third retrospective split, is what this cervical cancer study adds. The authors combined peritumoral ultrasound radiomics, transfer-learned deep features and clinical factors to predict lymphovascular space invasion before surgery, then carried the model through to disease-free survival and TCGA pathway analysis.

Key metrics: 292 women (188 training, 53 internal validation, 51 prospective test). Hybrid model AUC 0.79 (95% CI 0.66 to 0.92) with sensitivity 0.85 (0.71 to 0.98) in the prospective set, outperforming single-modality models. High-risk group showed shorter disease-free survival ($P<0.05$). SHAP identified the deep learning score as primary predictor; Wnt/beta-catenin signalling was enriched in LVSI-positive tumors.

Clinical relevance: Preoperative LVSI status informs surgical planning and risk stratification in cervical cancer, and ultrasound is the widely accessible modality for it. The honest reading: a 51-patient prospective test keeps the confidence interval wide, and the authors themselves call for larger multicenter validation before clinical application. The design worth copying is the prospective test set plus a biological plausibility check rather than another retrospective split.

PMID 42848195 DOI 10.1007/s00261-026-05804-7

Radiology research platform: automating multicenter radiological research workflows using a cloud-native platform.

Miltiadis I, et al. | [European radiology](#) | 2026-10-07

Before you budget another year of manual data wrangling for a multicenter imaging study, look at what a cloud-native architecture removes. This proof of concept connects institutional Orthanc DICOM edge nodes to a cloud backend through reverse-proxied REST, with a zero-footprint OHIF viewer and JSON-defined structured forms, so radiologists at multiple sites read assigned cases in one workspace without VPN setups or hardware shipments.

Key metrics: Proof-of-concept architecture description; the abstract reports no quantitative throughput or performance figures. Components: React client, Google Firebase cloud backend, localized Orthanc edge nodes keeping images on site, dynamic JSON evaluation forms, consensus or fixed-volume enrollment algorithms, embedded Open Health Imaging Foundation viewer.

Clinical relevance: Most departments fail multicenter validation not on statistics but on plumbing: moving images, forms and reader assignments across firewalls. Infrastructure of this shape lowers the cost of exactly the studies governance keeps demanding, held-out-site testing and multi-reader evaluations. For your AI committee this is a validation-capacity item with an EU AI Act Article 10 edge: when shared evaluation infrastructure is proposed, the architectural questions (where images stay, what metadata leaves, who controls the forms) are your data-governance documentation.

PMID [42844398](#) DOI [10.1007/s00330-026-12926-5](#)

An open vision-language model for diverse medical applications.

Sellergren A, et al. | [Nature medicine](#) | 2026-10-06

Open weights change who can adapt a medical foundation model. MedGemma is a collection of medical vision-language models built on Gemma 3, released with a medically tuned vision encoder (MedSigLIP) and designed for data-efficient fine-tuning across imaging domains including chest X-ray, while keeping the general capabilities of the base models.

Key metrics: Against similarly sized base models, on out-of-distribution tasks: 2.6 to 10% improvement in medical image question answering, 15.5 to 18.1% in chest X-ray finding classification, 10.8% in agentic evaluations. Fine-tuning MedGemma beat fine-tuning base Gemma 3 in limited-data settings; MedSigLIP performed comparably to or better than many specialized medical image encoders. Benchmark deltas only: the abstract reports no reader study and no clinical outcomes.

Clinical relevance: The most consequential foundation-model release in medical imaging this week, because the weights are open and the encoder is separable: research groups and vendors can fine-tune locally instead of calling a closed API. The published numbers are benchmark improvements, not clinical performance; what a fine-tuned derivative does on a specific casemix is an empirical question the paper explicitly leaves to downstream developers.

PMID [42839104](#) DOI [10.1038/s41591-026-04626-w](#)

The medical physicist's role in the artificial intelligence era: an updated EFOMP curricular and professional framework.

Kagadis GC, et al. | [Physica medica](#) | 2026-10-05

Decide who in your department owns AI validation competence before the next tool arrives; EFOMP has now written the map. Five years after its first AI curriculum white paper, the federation reports that the radiology, radiotherapy and nuclear medicine core curricula all carry AI as a core competency, and proposes an advanced cross-subspecialty framework in four domains: validation and quality assurance, AI lifecycle management, generative AI and foundation models, and governance.

Key metrics: Consensus and curriculum-analysis document, no performance figures. Inputs: four primary source documents, three large-scale surveys of medical physicists (2020, 2024, 2025), structured iterative expert deliberation. Findings: the three updated curricula diverge in ECTS allocation, coverage of recent AI architectures and treatment of the EU AI Act; six competency gaps common to all three subspecialties.

Clinical relevance: The first professional framework translating EU AI Act responsibilities into a named clinical profession's competency ladder, covering how competence is acquired in training, maintained through accredited CPD and formally recognized across a career. For your AI committee this is an Article 4 literacy item: a defensible template for documenting who is competent to perform acceptance testing and lifecycle monitoring, and the six named gaps double as an audit checklist for your own training plan.

PMID [42833069](#) DOI [10.1016/j.ejmp.2026.107209](#)

An Inline-Integrable Image Restoration Model With ROI Fine-Tuning and Slice-to-Volume Reconstruction for Accelerated Whole-Brain Vessel Wall Imaging.

Wang P, et al. | *Magnetic resonance in medicine* | 2026-10-08

Sequence acceleration is the rare AI application whose benefit shows up on the schedule the same week. This transformer-based 2.5D restoration model brings whole-brain vessel wall MRI under 6 minutes on standard 20-channel coils by restoring CAIPIRINHA 2x2-accelerated acquisitions to GRAPPA twofold quality, with region-of-interest fine-tuning and slice-to-volume reconstruction, built to run inside the vendor's inline reconstruction pipeline.

Key metrics: Peak SNR improved by more than 7 dB and SSIM by 0.19 over input images in the retrospective test set. Applied to fourfold-accelerated data, significant gains in vessel wall SNR, wall-to-lumen and wall-to-CSF CNR and image quality scores versus offline or inline GRAPPA. Improvements held in an independent multi-center cohort and reproduced on a smaller field of view with a 64-channel coil. Siemens 3 T; prospective multi-center testing included.

Clinical relevance: Vessel wall imaging is clinically valuable and chronically too slow; a 6-minute whole-brain protocol on ordinary hardware moves it from research luxury toward routine stroke and vasculopathy workup. The design detail worth copying: generalization was tested across sites, coils and field-of-view settings, and the model was engineered for inline deployment rather than offline demonstration.

PMID 42850761 DOI 10.1002/mrm.70594

Decoupling Multi-Contrast Super-Resolution: Self-Supervised Implicit Re-Representation for Unpaired Cross-Modal Synthesis.

Wu Y, et al. | *IEEE transactions on medical imaging* | 2026-10-09

Paired training data is the standing obstacle to multi-contrast MRI super-resolution; this work removes it by decoupling the problem into a population-level unpaired cross-modal synthesis prior, trained once, and a lightweight patient-specific implicit re-representation module optimized self-supervised at inference.

Key metrics: The abstract makes comparative claims without numeric values: superior quantitative performance across datasets and robustness at 16x and 32x upsampling, a regime where competing methods are reported to fail. Scale-agnostic by construction; no paired low/high-resolution or paired cross-modal training data required, though inference still uses a matched high-resolution reference contrast from the same patient.

Clinical relevance: Dropping the paired-data requirement matters because paired acquisitions are what make MRI enhancement models expensive to train and brittle across scanners and protocols. A methods contribution at this stage: super-resolved images feed downstream measurement, so clinical adoption will need task-based evaluation on real cohorts rather than image-similarity scores alone.

PMID 42853735 DOI 10.1109/TMI.2026.3739662

UPER: A Unified Polyhedral Representation for universal multi-dimensional diffusion MRI image modeling.

Yao T, et al. | *Medical image analysis* | 2026-10-06

Diffusion MRI has resisted foundation-model treatment because every acquisition scheme differs; UPER recasts multi-dimensional dMRI into a unified polyhedral representation and pretrains a masked autoencoder across 2,800 scans from 17 datasets spanning more than 10 unique shell configurations, in one end-to-end transformer framework.

Key metrics: 2,800 scans, 17 datasets, over 10 unique shells; masked autoencoder with self-supervised plus supervised training. The abstract claims superior performance on various downstream tasks versus traditional deep learning frameworks without reporting specific values. Code publicly available on GitHub.

Clinical relevance: The first deep learning model claimed to handle varying radial and angular sampling and heterogeneous shell configurations in a single framework, the precondition for dMRI models that transfer across studies instead of being rebuilt per protocol. For neuroimaging groups the practical gain is a pretrained starting point for microstructure and connectivity work on whatever acquisition they already have.

PMID 42854624 DOI 10.1016/j.media.2026.104336

SECTION 2

INDUSTRY & REGULATION

*Sources: Official registers | Press releases | FDA databases***[FDA] AIRS Medical's SwiftSight Body Composition cleared for quantitative MRI body composition***AIRS Medical | 5 October 2026*

Cleared for what: 510(k) clearance for software that quantifies muscle and adipose tissue from MR images, including visceral and subcutaneous fat, abdominal and thigh muscle, muscle quality and bilateral balance, plus quantitative liver fat, benchmarked against age- and sex-matched reference populations. On what evidence: substantial equivalence to existing quantitative MR body-composition software; no product-specific accuracy study has been published, and the underlying liver-fat figures cited in coverage describe the PDFF reference technique, not this product. First deployment is SimonMed Imaging's self-pay longevity screening service. Questions for the vendor: the predicate and the equivalence basis; product-specific accuracy and repeatability against a dedicated PDFF protocol on your scanners; the composition of the normative reference database by scanner and population; and how an incidental liver-fat flag from a self-pay screen is expected to enter a billed diagnostic pathway. Monitoring regardless of clearance: sample concordance of its liver-fat values against your reference method before any report leaves the department.

Source: [PR Newswire](#)**[FDA] MRI2CT's NextMR cleared for 3D lumbar spine bone visualization from standard MRI***MRI2CT | 7 October 2026*

Cleared for what: 510(k) clearance for generating 3D bone visualization of the lumbar spine and sacrum from routine 1.5 T and 3 T diagnostic MRI sequences, with output in under 60 seconds through a cloud application; the company's own release footnotes that NextMR images are not intended to replace CT in general. On what evidence: a vendor-reported retrospective validation on 145 adults claiming submillimeter cortical delineation accuracy across patient subgroups; no peer-reviewed publication is cited. Questions for the vendor: where the 145-patient validation is published and what the acceptance threshold was; performance in the cases that matter for surgery (spinal hardware, deformity, transitional anatomy); the defined failure modes and exclusion criteria; whether surgical-planning use still requires CT confirmation; and where the cloud processing holds your data. If adopted, log every case in which a surgeon requests a confirmatory CT anyway: that rate is your real-world utility metric.

Source: [SPINEMarketGroup](#)**[REGULATION] Avatar Medical Vision CE-marked for glasses-free 3D visualization with Barco displays***Avatar Medical | 8 October 2026*

Cleared for what: CE marking, announced at JFR 2026, for software that visualizes CT and MR images in 3D and VR for surgical planning and intervention support, now certified for use with Barco's Eonis glasses-free 3D medical display; it renders patient-specific anatomy without segmentation. This is a visualization claim, not a detection or diagnostic-performance claim, and nothing in the release reports reader or outcome evidence. On what evidence: conformity assessment building on the product's initial CE marking of March 2025; the news is the certified pairing with a specific display. Questions for the vendor: the exact certified intended-use sentence and risk class; whether the 3D rendering is an adjunct to, not a substitute for, the standard 2D read; the display calibration and QA requirements the certification assumes; and which specialties the claimed workflows were validated with.

Source: [Avatar Medical](#)

[MARKET] Bayer sells Blackford Analysis to Sopris Capital; platform to combine with Azra AI*Sopris Capital / Blackford | 9 October 2026*

What happened: Sopris Capital completed the acquisition of AI-orchestration vendor Blackford Analysis from Bayer and is folding it under its portfolio company Azra AI; Blackford continues as a separate legal entity and, per the release, existing deployments and operations continue unchanged. Why it matters to a buyer: the orchestration layer that routes your imaging AI has now changed hands twice in under four years (independent, Bayer, Sopris), and its 130-plus applications from 45-plus vendors are contracts that follow the owner. Evidence level: company announcement only. Questions for your account manager: which contractual commitments (SLAs, data processing agreements, roadmap, exit assistance) survive the change of control; what the combination with an oncology-workflow company means for radiology R&D priority; and what happens to per-tool validation records and version history through the migration. Monitoring regardless: keep your own inventory of which algorithms run through the platform, with versions and dates, so a platform transition never erases your audit trail.

Source: [PR Newswire](#)**[PARTNERSHIP] CARPL.ai to distribute HOPPR foundation models for local fine-tuning***CARPL.ai / HOPPR | 7 October 2026*

What happened: CARPL.ai announced a partnership giving healthcare providers access to HOPPR's imaging foundation models to develop, fine-tune and deploy custom radiology AI applications, such as report generation tuned to institutional style, through the CARPL platform. Why it matters: this moves foundation-model fine-tuning from research groups to hospital procurement, the same shift MedGemma's open weights make on the open-source side this week. Evidence level: partnership announcement; no cleared product is named. Questions before piloting: who is the legal manufacturer of a model your team fine-tunes on CARPL, and who carries the regulatory obligations that follow; what validation and monitoring tooling is included for models you build; what happens to your local training data and the resulting model weights if the partnership or your contract ends; and whether any fine-tuned output class (report drafts above all) would be a regulated device in your jurisdiction before clearance exists for it.

Source: [HOPPR](#)

SECTION 3

POST-MARKET SIGNAL*Drift | Recalls and MAUDE | Real-world monitoring | EU AI Act Article 72***Digital Omnibus aftermath: the Article 72 monitoring template becomes voluntary guidance due September 2027***9 October 2026*

Two compliance analyses published this week work through what Regulation (EU) 2026/1744 (the Digital Omnibus on AI, consolidated text since 27 July 2026) did to post-market monitoring under the EU AI Act. The implementing act that was to fix a mandatory monitoring-plan template by 2 February 2026 was never adopted and has been replaced: the Commission now owes guidance, including a template the amending act's recitals call voluntary, by 2 September 2027. Article 72's substantive duty (actively and systematically collect and analyse lifetime performance data, with the plan in the Annex IV technical file) is unchanged and formally applies from 2 August 2026, while the high-risk classification rules it attaches to now apply from 2 December 2027 (Annex III) and 2 August 2028 (Annex I, which covers MDR devices); the analyses flag the fit between those dates as an open question. The practical consequence for imaging: anyone waiting for an official template before starting monitoring is waiting for a voluntary document arriving late next year. Write the plan against Article 72(1) and (2) now (metrics, baselines, data sources, thresholds, review cadence) and integrate it with the MDR post-market surveillance file your vendors already owe you.

Link: <https://www.complipath.io/guides/eu-ai-act-post-market-monitoring/>

Medical-device reading of the amended AI Act: monitor AI-to-AI interactions, build logging in

9 October 2026

A device-regulatory analysis published 9 October maps the amended AI Act onto medical-device quality systems: MDCG 2025-6 identifies monitoring of interactions with other AI systems, devices and software as the main genuinely new post-market expectation, and Article 12 logging is a design requirement that must exist in the product, not only in a procedure, because it feeds traceability and post-market monitoring. For a radiology department whose reading chain now stacks triage, detection, reconstruction and reporting tools from different vendors, interaction monitoring is not an abstraction: it is whether you can tell which model saw what when two of them disagree. The December 2025 proposal to move the MDR from Annex I Section A remains at committee stage, so the current integration rules apply.

Link: <https://candelalabs.tech/blog/ai-act-medical-devices-digital-omnibus>

SECTION 4

PLAYBOOK SNIPPET

One concrete step your AI committee can take this month

EU AI Act Article 4 (AI literacy), applying since February 2025. The AI Act application calendar is under revision following the Digital Omnibus proposal, so verify enforcement dates against the current text

Step: Pick one deployed model. Produce a single signed and dated page: what it is cleared for, what your acceptance test showed, which patient subgroups and reader groups were checked, and what triggers suspension.

Why: A certificate proves attendance; a completed audit page proves the person can tell when the model is failing in their own population, which is what an inspector, and your own quality committee, will actually ask.

Worked example: Sun Y et al. (Academic Radiology, PMID 42855384, DOI 10.1016/j.acra.2026.09.036) show what the reader-group line of that page should contain: external-cohort AUC 0.939, senior readers improved from 0.868 to 0.929 with assistance, juniors from 0.548 to 0.659. If your acceptance test cannot show this split for your own readers, the audit page is not complete.

SECTION 5

MEDIA HIGHLIGHTS

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

AI cuts turnaround times in high-volume imaging, but the pipes decide

AuntMinnie.com | 8 October 2026

AuntMinnie covered a JACR study of a 4.5-year, 20-center multi-vendor AI implementation (10 tools, 7 vendors, 96,874 exams): AI availability was associated with 26% lower median report turnaround for trauma radiography and 18% for knee MRI, at an annualized program cost of roughly a quarter of one radiologist FTE. The governance finding inside the efficiency story: 72% of AI latency came from data routing, and 7.2% of AI results arrived only after the report was signed, 13.2% for chest CT. Turnaround gains and adoption (91.4% active use) count for little in the cases where the result misses the reporting window; per-modality latency against report finalization is the line to add to your monitoring.

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

JFR 2026: deployment, drift and the usage gap dominate the emergency-AI agenda

AuntMinnie Europe | 8 October 2026

From the opening day of JFR 2026 in Paris: the dedicated emergency-imaging AI session argued that a published performance result never guarantees stable local performance, making continuous monitoring including drift surveillance part of deployment, not an afterthought. The most sobering number quoted: in a Lancet-published randomized evaluation of an AI-assisted stethoscope across roughly 700,000 patients, 40% of practitioners had stopped using the tool within twelve months, and actual usage rate, not algorithmic accuracy, determined population impact. Emergency-specific tools account for more than half of the imaging-AI market, so the adoption question lands squarely on radiology.

Link: <https://www.auntminnieeurope.com/imaging-informatics/artificial-intelligence/art...>

SwiftSight analysis: a 510(k) clearance arriving through a self-pay screening channel

Regulatory News | 7 October 2026

An in-depth analysis of the week's AIRS Medical clearance makes the buyer-relevant distinctions the press release does not: the product was cleared on substantial equivalence with no published product-specific accuracy trial; the sensitivity and specificity figures circulating describe the underlying PDFF biomarker, not the product; and the first deployment channel is a nationwide self-pay longevity MRI program, so reports will reach referring physicians without a payer or guideline gate. New Category III CPT codes (0697T, 0698T) exist but carry no fixed Medicare fee. A clean worked example of why the intended-use sentence and the evidence level, not the clearance itself, should drive what your clinicians do with the number.

Link: <https://regulatorynews.com/swiftsight-body-composition-liver-fat-clearance>

SECTION 6

PROFESSIONAL SOCIETIES

SIIM | ACR | RCR | Official publications and event coverage

EuSoMII Annual Meeting 2026 held in Heraklion with post-market surveillance on the main program

9 to 10 October 2026

EuSoMII's Annual Meeting took place in Heraklion, Crete, this week, with a program that reads like a governance syllabus: a dedicated 'Post-Market Surveillance: A How to Guide' session, an ACR Data Science Institute presentation of its AI governance and monitoring programs including the Assess-AI registry, joint sessions with ESNR, SIIM, EFRS and MICCAI on safe integration of AI tools and human-AI interaction paradigms, and a keynote on readiness for the European Health Data Space.

Source: <https://www.eusomii-annualmeeting2026.com/schedule>

JFR 2026 in Paris puts practical AI integration at the center of the French national congress

8 to 10 October 2026

The Journées Francophones de Radiologie, the French-speaking world's national radiology congress, ran in Paris this week with AI prominent across the scientific program, from a dedicated session on integrating AI into emergency imaging (chest, abdomen, neuroradiology, triage and decision support) to vendor portfolio launches timed to the meeting. The session's through-line matched this digest's: published performance does not guarantee local performance, and monitoring for drift belongs to deployment.

Source: <https://www.auntminnieeurope.com/imaging-informatics/artificial-intelligence/article/15837079/jfr-how-to-integrate-ai-in-emergency-imaging>

RSNA introduces an AI certificate program for imaging professionals

4 October 2026

RSNA announced an artificial intelligence certificate program aimed at giving imaging professionals the knowledge and skills to assess and use AI technologies in their current practice. In an EU department, structured certification maps directly onto the Article 4 AI-literacy obligation, but the playbook distinction stands: a certificate documents training, while a completed model audit documents competence. Use programs like this as the training layer under a named, auditable owner for each deployed model.

Source: <https://www.appliedradiology.com/news/rsna-launches-ai-certificate-program-for-imaging-professionals>

SECTION 7

KEY OPINION LEADERS

*LinkedIn | Public statements | Week of 4 to 10 October 2026***Daniel Pinto dos Santos***Radiologist and imaging informatics researcher; EuSoMII*

Posted 9 October from the EuSoMII Annual Meeting, amplifying the ACR Data Science Institute's Assess-AI message that 'the most useful evidence about clinical AI comes after deployment' after Woojin Kim's presentation: 'Really nice presentation by Woojin Kim at #eusomii2026!! Definitely something urgently needed towards post-market surveillance in the EU as well.' A European informatics voice explicitly calling for an EU equivalent of the ACR's post-market monitoring registry.

Source: [LinkedIn](#)**Dr Amine Korchi***Radiologist and HealthTech investor, Geneva*

Posted 9 October, framing his new personal investment with the question his followers kept asking after Gleamer's exit: 'So, what's the next Gleamer?' His answer is not an imaging company but Vocca, AI voice agents that handle clinic phone calls end to end; he cites growth from 2,000 to 15,000 practitioners in a year, 1,500 practices and over one million patient interactions a month, with 70% of conversations resolved without human involvement. A signal of where radiology-adjacent AI money is moving: from detection algorithms to operational agents.

Source: [LinkedIn](#)**Matt Lungren***Chief Data and AI Officer, Oku Health; formerly Microsoft Health AI*

Posted 7 October announcing why he left Microsoft in January: the launch of Oku Digital Health from Topcon Healthcare, positioning routine eye exams as 'a front door for early disease detection and accelerated clinical development', with Lungren as Chief Data and AI Officer. One of imaging AI's most visible figures betting that the next clinical AI platform grows out of ophthalmic imaging rather than radiology.

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 4 to 10 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.