DR. SERGEY MOROZOV

WEEKLY AI INTELLIGENCE REPORT

Radiology and Medical Imaging

Week 34 | 16 to 22 August 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	5	2	3
Peer-reviewed papers	Industry / FDA items	Media highlights	Key opinion leaders

GOVERNANCE TAKEAWAY OF THE WEEK

If chest radiography AI is live or planned in your department, schedule a local multi-reader check this month and score accuracy per finding type, not pooled. In a prospective crossover study of four commercial tools, readers became faster and more confident while accuracy stayed flat and fell for effusions and nodules through added false positives. Speed and confidence are exactly what a demo shows you; neither is evidence of accuracy. The same week's osteonecrosis meta-analysis measured the discount to apply to internal validation claims: external testing roughly halved the diagnostic odds ratio.

SECTION 0

EXECUTIVE SUMMARY

Key trends, week of 16 to 22 August 2026

[FLAGSHIP] LIVER AI CLEARS A PRE-SPECIFIED TRIAL ENDPOINT INSIDE A ROUTINE WORKFLOW

Zhang X et al. (Nature Medicine) validated the LiON liver malignancy system on 22,251 patients (AUC 0.975), then ran a registered single-arm trial in 10,333 patients with LiON as an additional reader: primary endpoint met at AUC 0.952, 51 overlooked lesions surfaced, 37 reports amended, 22 multidisciplinary escalations. Subgroups were reported where liver imaging is hardest: cirrhosis at AUC 0.924.

[READER STUDY] FOUR COMMERCIAL CHEST X-RAY AIS: FASTER, MORE CONFIDENT, NOT MORE ACCURATE

Lemke T et al. (Academic Radiology) tested four commercial chest radiography tools with five readers over 1200 consecutive patients in a prospective crossover design. Accuracy did not improve and fell for effusions and nodules through false positives, while reading got 6 to 17 s faster for most residents and confidence rose. The gap between felt benefit and measured benefit is the automation bias signature.

[VALIDATION GAP] EXTERNAL VALIDATION ROUGHLY HALVES THE DIAGNOSTIC ODDS RATIO

Lu F et al. (J Med Internet Res) pooled 12 osteonecrosis-of-the-femoral-head AI studies, 16,189 hips: sensitivity 0.91, specificity 0.95 overall, but externally validated models carried a DOR of 129 against 230 for internal-only validation. The discount between a vendor's internal figures and external reality now has a measured size.

[EXPERIENCE GAP] PITUITARY MICROADENOMA ASSISTANCE MOVES EVERY READER, SENIORS INCLUDED

Kang S et al. (Neuroradiology) combined clinical, radiomics and deep learning features on non-contrast T1COR MRI and ran a three-round study with five radiologists over 315 lesions: accuracy gains at every experience level, up to roughly 30 absolute percentage points in the hardest external cohort, with improved inter-reader consensus.

[WATCHPOINT] FELT BENEFIT IS NOT MEASURED BENEFIT

Three results this week put numbers on the gap between what performance claims feel like and what local testing finds: chest x-ray readers got faster and more confident with no accuracy gain, external validation halved the diagnostic odds ratio in the osteonecrosis meta-analysis, and a meningioma model whose tumor Dice survived external testing saw its edema output collapse from 0.63 to 0.31. The common instruction: test the claim you will actually rely on, in your own population, before it shapes reader behavior.

SECTION 1

PEER-REVIEWED PAPERS

Source: PubMed | Verified indexing | No preprints | 16 to 22 August 2026

Large-scale AI-guided liver malignancy diagnosis: multicenter study and a single-arm trial.

Zhang X, et al. | ***Nature medicine*** | *2026-08-19*

If you want the current reference standard for deployment evidence in imaging AI, read this one. LiON, a contrast-enhanced CT system for liver malignancy diagnosis, was trained on 6,443 patients, retrospectively validated across 22,251, then tested in a prospective single-arm trial in 10,333 patients as an additional AI reader inside the existing workflow, against a pre-specified primary endpoint. The authors are explicit that comparative prospective evidence on clinical outcomes is still missing.

Key metrics: Training 6,443 patients; multicenter and real-world validation 22,251 patients, AUC 0.975 (95% CI 0.971-0.979); hepatic steatosis AUC 0.971 (0.952-0.985), cirrhosis AUC 0.924 (0.901-0.946). Single-arm trial in 10,333 patients: primary endpoint met, AUC 0.952 (0.942-0.961), lower CI bound above 0.900. AI-human collaboration surfaced 51 previously overlooked lesions (15 malignancies), 37 amended reports and 22 multidisciplinary team escalations. NCT07153783.

Clinical relevance: Kept on scientific merit. This is what a diagnostic safety net looks like when it is evaluated seriously: a registered trial with a pre-specified endpoint, deployment as an additional reader rather than a replacement, and subgroup performance reported exactly where liver imaging gets hard, in steatosis and cirrhosis, instead of pooled away. The measured output was not only discrimination but amended reports and escalations that changed management.

PMID 42618635 **DOI** 10.1038/s41591-026-04589-y

Artificial Intelligence-Assisted Chest Radiography: A Prospective Crossover Multi-Reader Study on Diagnostic Performance and Workflow Efficiency.

Lemke T, et al. | ***Academic radiology*** | *2026-08-21*

Do not assume a cleared chest radiography AI improves your readers; measure it before it shapes their behavior. Five readers assessed 1200 consecutive patients (1861 radiographs) for five finding types without AI and with each of four commercially available algorithms, in a prospective crossover design with randomized case order and a 14-day washout. AI assistance did not improve diagnostic accuracy, and for pleural effusions and pulmonary nodules accuracy fell in several reader-tool pairings through added false positives.

Key metrics: 1200 patients, 1861 radiographs, five readers with one to six years of experience, four commercial AI tools, five target findings. Accuracy: no overall improvement; decreases for pleural effusions and pulmonary nodules in several pairings due to false positives. Interpretation time: three of five residents faster, median reductions of 6 to 17 s per case ($p<=0.031$). Four of five readers reported higher diagnostic confidence; senior consultations fell in selected pairings, CT escalation fell in one.

Clinical relevance: The dissociation is the finding: readers became faster and more confident while accuracy stayed flat or fell, which is the signature of automation bias. For your AI committee this changes acceptance testing: a local multi-reader assessment per finding type, with the false positive rate tracked separately, belongs in the go-live file, and speed or confidence must not be accepted as surrogate endpoints for accuracy (EU AI Act Article 14 human oversight, Article 15 accuracy and robustness).

PMID 42629289 **DOI** 10.1016/j.acra.2026.08.014

Diagnostic Accuracy of Medical Imaging-Based Artificial Intelligence for Osteonecrosis of the Femoral Head: Systematic Review and Meta-Analysis.

Lu F, et al. | ***Journal of medical Internet research*** | *2026-08-20*

Before you quote internally validated performance for femoral head osteonecrosis AI, apply the discount this meta-analysis measured. Across 12 studies and 16,189 hip joints, pooled diagnostic performance was high, but models with only internal validation carried a diagnostic odds ratio nearly twice that of externally validated ones.

Key metrics: 12 studies, 16,189 hip joints. Pooled sensitivity 0.91 (95% CI 0.87-0.95), specificity 0.95 (0.93-0.96), SROC AUC 0.97 (0.95-0.98); heterogeneity I2 = 72%. MRI-based models: DOR 382 (220-665) versus x-ray 106 (60-190). External validation DOR 129 (51-329) versus internal-only 230 (104-510). PROSPERO CRD420261307216.

Clinical relevance: For your AI committee this changes validation: treat internal-only figures as an upper bound and require external, ideally local, performance before acceptance (EU AI Act Article 15 accuracy and robustness). The internal-to-external gap measured here, roughly a halving of the diagnostic odds ratio, is a usable default discount when a vendor dossier offers nothing better.

PMID 42623169 **DOI** 10.2196/95648

Multicenter External Validation of an AI-Based Funduscopic Carotid Atherosclerosis Score and Assessment of Its Association With Coronary Artery Calcification: External Validation Study.

Han C, et al. | ***JMIR medical informatics*** | *2026-08-19*

External validation at scale, with honest numbers: an AI-derived fundus score for carotid atherosclerosis held up across 108,982 health checkup participants at five Korean centers, at deliberately modest discrimination. The claim worth reading is not the AUC but the independent association with coronary artery calcification after adjusting for the pooled cohort equations.

Key metrics: 108,982 participants, five health promotion centers, 2018 to 2021. AUC for carotid atherosclerosis 0.700 (95% CI 0.697-0.703). Per 10% absolute increase in DL-FAS, adjusted for the PCE 10-year risk score: odds ratio 1.29 (1.28-1.31) for carotid atherosclerosis and 1.29 (1.24-1.33) for coronary artery calcification; associations held under age 60 and in PCE low- and moderate-risk subgroups.

Clinical relevance: As opportunistic screening evidence goes, this is the right shape: a multicenter external cohort, adjustment against the risk score already in clinical use rather than against nothing, and discrimination reported without inflation. An AUC of 0.700 will not gate anyone into angiography, but for a zero-added-acquisition signal from images already taken, incremental association beyond the existing risk equations is the relevant claim, and it survived in the subgroups where early intervention matters most.

PMID 42616614 **DOI** 10.2196/77335

Automated deep learning-based segmentation and volumetric analysis of meningiomas.

de Wilde D, et al. | ***Neuroradiology*** | *2026-08-20*

A cautionary pair of numbers for anyone considering automated meningioma volumetry: tumor segmentation generalized to an external dataset almost unchanged, while peritumoral edema segmentation collapsed. Both outputs come from the same model, trained on 100 patients and externally tested on 88.

Key metrics: nnU-Net trained on contrast-enhanced T1 and FLAIR from 100 University Hospital Zurich patients; external validation on 88 cases from the TCIA meningioma SEG-Class dataset. Dice for meningioma: 0.87 +/- 0.23 internal, 0.86 +/- 0.17 external. Dice for peritumoral edema: 0.63 +/- 0.38 internal, 0.31 +/- 0.35 external.

Clinical relevance: The tumor result supports automated longitudinal volumetry of the lesion itself, the task that consumes the most reader time in meningioma surveillance. The edema result shows how a co-packaged secondary output can fail external testing outright while the headline number survives; the authors themselves call for larger multicenter validation before clinical use.

PMID 42622848 **DOI** 10.1007/s00234-026-04152-z

A stacking model for AI-assisted diagnosis of suspected pituitary microadenomas on non-contrast T1COR MRI: a multicenter reader study on bridging the experience gap.

Kang S, et al. | ***Neuroradiology*** | *2026-08-20*

If AI assistance is supposed to bridge the experience gap, test it against the gap you actually have. A stacking model combining clinical, radiomics and deep learning features on non-contrast T1COR MRI improved diagnostic accuracy for suspected pituitary microadenoma across all experience levels in a three-round study of five radiologists, with the largest gains in the hardest external cohort.

Key metrics: 636 patients from three centers; training n = 321, internal validation n = 138, external cohorts n = 136 and n = 41. Combined model AUC 0.818 (EVC1) and 0.899 (EVC2); sensitivity 0.929, specificity 0.833 in EVC2. Reader study: 315 lesions, five radiologists (three junior, two senior). With combined-model assistance seniors reached 75.4 to 79.0% accuracy internally (adjusted $p < 0.05$) and up to 85.4% in the most challenging external cohort, an absolute gain of approximately 30% over unaided reading; inter-reader consensus improved in both lesion groups.

Clinical relevance: For your AI committee this changes acceptance testing: measure assistance effects per experience stratum, not pooled. A tool that moves juniors and seniors by different amounts changes your supervision and double-reading arrangements, which is an EU AI Act Article 14 human oversight question, not only an accuracy one.

PMID 42622847 **DOI** 10.1007/s00234-026-04135-0

A comprehensive framework for multi-class prostate cancer classification using biparametric MRI: a multi-center retrospective study.

Zheng B, et al. | ***Insights into imaging*** | *2026-08-17*

Multiclass risk stratification, not a binary cancer call, is the clinically useful question before prostate biopsy, and this framework answers it end to end: automated prostate and lesion segmentation, feature extraction and three-way ISUP-based classification on biparametric MRI across 2543 patients.

Key metrics: 2543 patients with suspected clinically significant prostate cancer. Segmentation: 0.971 internal validation, 0.934 external testing. Integrated classifier combining radiomics, learned deep features and clinical variables: macro-AUC 0.87 (0.84-0.90) across internal and external evaluation, outperforming total PSA, PSA density and PI-RADS, with robustness across most PSA and PI-RADS 4 subgroups.

Clinical relevance: The step past binary csPCa detection matters clinically: separating ISUP 2-3 from ISUP 4-5 before biopsy is what changes management, from surveillance eligibility to treatment planning. The comparison set was the right one, the clinical variables a urologist would otherwise use, and the external macro-AUC held at the pooled figure.

PMID 42606778 **DOI** 10.1186/s13244-026-02367-5

Motion Artifact-Aware Self-Supervised Representation Learning for 3D Brain MRI Motion Artifact Reduction.

Safari M, et al. | ***IEEE transactions on medical imaging*** | *2026-08-17*

Motion correction without paired data is the constraint that decides whether a method can leave the lab, and this framework is built inside it: self-supervised representation learning for 3D brain MRI motion artifact reduction that needs neither clean-corrupted pairs nor k-space access nor motion labels.

Key metrics: In-silico evaluation: PSNR 23.81 dB, SSIM 91.55%, NMSE 0.79%. On the in-vivo MR-ART dataset, up to 2.0 dB PSNR gain over a source-only supervised model after unsupervised domain adaptation, remaining within 0.25 to 0.47 dB of an oracle supervised model trained on real paired data. Volumetric error in the corpus callosum and ventricular system fell by more than 50% at the milder motion level.

Clinical relevance: The design constraint is the contribution: paired clean and corrupted acquisitions rarely exist in clinical archives, so a method that trains without them and adapts to a new site without labels is the kind that can plausibly reach deployment. The gains were demonstrated on structural quantification, exactly where motion silently biases volumetric studies at scale.

PMID 42606966 **DOI** 10.1109/TMI.2026.3724112

SECTION 2

INDUSTRY & REGULATION

Sources: Official registers | Press releases | FDA databases

[FDA] Cortechs.ai announces FDA clearance for NeuroQuant PET

Cortechs.ai | 18 August 2026

Cleared for: automated quantification, regional analysis and structured reporting of PET performed for the assessment of cognitive impairment and other neurological conditions (510(k) K261916), generating SUVRs and Centiloid values for the amyloid tracers flutemetamol, florbetaben and florbetapir. Evidence level: vendor materials; the announcement cites no reader study or clinical performance data. Questions before buying: which atlas and normative database drive the quantification, what is test-retest variability of SUVR and Centiloid outputs on your scanners, how does accuracy differ between PET-CT and PET-MR workflows, and what happens outside the three named tracers. Regardless of clearance, run a local concordance check of Centiloid output against your current visual reads before results feed dementia care pathways.

Source: PR Newswire

[FDA] FDA clears CureMetrix cmAngio 2.0 with BAC quantification and age-based context

CureMetrix | 17 August 2026

Cleared for: detection and now quantification of breast arterial calcification on screening mammograms, with age-based percentile and prevalence context from a reference library of more than 170,000 studies. Evidence level: vendor validation testing, reported AUC 0.98 to 0.99 across FFDM and DBT on GE HealthCare and Hologic systems in women aged 40 to 90; no peer-reviewed outcome data in the announcement. Questions: was the AUC measured at screening prevalence or on an enriched set, how repeatable is the quantification across vendors and compression, who acts on a high BAC percentile and through which cardiology pathway, and what does the tool add over radiologist BAC reporting now that the 2025 BI-RADS Atlas recognizes the finding. Note the downstream obligation: Maryland-style patient notification laws make BAC a reportable result, so define the communication workflow before go-live.

Source: EIN Presswire

[REGULATION] FDA opens comment period on regulating generative AI-enabled medical devices

FDA CDRH | 18 August 2026

Not a clearance: CDRH published a discussion paper (docket FDA-2026-N-7874, comments due 19 October 2026) outlining how it might regulate generative AI-enabled devices: a two-axis risk framework, competency-based premarket evaluation combining device benchmarking with clinical confirmation of the final user-facing device, and risk-proportionate postmarket monitoring. The consequential idea for deployers: the agency is weighing whether to accept greater premarket uncertainty in exchange for heavier reliance on postmarket monitoring, which would shift validation burden toward the institutions that deploy these tools. If you run or plan generative reporting tools, comment on the docket; the evidence bar negotiated here is the one your future vendors will be graded against.

Source: MedTech Dive

[PARTNERSHIP] Azra AI and 4DMedical partner on enterprise lung health intelligence

Azra AI / 4DMedical | 17 August 2026

What it is: an OEM and reseller agreement, not a clearance. 4DMedical will resell Azra AI's report-mining patient identification and care coordination platform and integrate its own quantitative lung imaging analytics, including the FDA-cleared CT:VQ, across incidental-finding and lung health pathways. Evidence level: vendor materials; the specific product integrations are described as planned and no combined-workflow performance data are cited. Questions: which findings does the report NLP surface with what sensitivity and false positive rate on your local report style, who owns follow-up once the platform routes a patient, and does the integrated offering require new regulatory review as the intended use expands from patient identification toward severity assessment.

Source: PR Newswire

[MARKET] Nanox and Vertec Scientific sign exclusive UK reseller agreement for HealthOST

Nanox AI | 18 August 2026

What it is: an exclusive three-year UK reseller agreement with minimum annual license commitments for Nanox.AI's HealthOST bone solution, which flags patients at risk of osteoporosis and vertebral fractures from routine CT; a distribution deal, not a new clearance and not new evidence. Questions for UK buyers: which HealthOST outputs fall inside the UK intended use, what published performance exists for a UK population and scanner mix, and how opportunistic findings will enter a funded fracture liaison pathway. Opportunistic screening tools only return value where the downstream service exists; without a commissioned pathway the flag is a liability, not a benefit.

Source: GlobeNewswire

SECTION 3

POST-MARKET SIGNAL

Drift | Recalls and MAUDE | Real-world monitoring | EU AI Act Article 72

Software issues prompt Class 2 recall of nearly 11,000 Philips Azurion interventional systems

Radiology Business | 18 August 2026

An FDA Class 2 recall covering the Philips Azurion R2.2.10 interventional fluoroscopy platform was reported this week: two software issues can degrade image quality or disable geometry movements, with potential for delayed therapy, procedural complications or inappropriate treatment in exactly the patients least able to tolerate delay. Philips sent Urgent Medical Device Correction letters in mid-July with interim workarounds, a cold-restart procedure pending a software update; nearly 11,000 units worldwide are affected. The post-market lesson for imaging departments: software regression in a mature, widely deployed platform is a live failure mode, and your incident path from a technologist noticing artifacts to the vendor complaint file is the mechanism that makes a recall like this visible. Verify that path exists and is used.

Link: https://radiologybusiness.com/topics/healthcare-management/healthcare-policy/sof...

SECTION 4

PLAYBOOK SNIPPET

One concrete step your AI committee can take this month

EU AI Act Article 14 (human oversight)

Step: Before deployment, record each reader's baseline on a fixed local case set: reading time and detection or marking rate per finding type. Repeat at 3 months with the tool on, and report per-reader deltas; pooled averages hide the effect.

Why: Human oversight is only meaningful if you know what the tool does to each specific human. A tool that speeds up one reader and degrades another looks neutral in the average, and the readers themselves cannot feel the difference: confidence rises whether or not accuracy does.

Worked example: Lemke T et al. tested four commercial chest radiography AIs with five readers over 1200 consecutive patients: accuracy fell for effusions and nodules in specific reader-tool pairings while three of five residents got faster and four of five felt more confident. Only the per-pairing analysis surfaced this. Academic Radiology, PMID 42629289, DOI 10.1016/j.acra.2026.08.014.

SECTION 5

MEDIA HIGHLIGHTS

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

Radiology Partners petitions FDA for greater clarity around imaging AI regulations

Radiology Business | *19 August 2026*

Radiology Partners' technology division Mosaic Clinical Technologies filed an FDA citizen petition (submitted 12 August, reported this week) asking the agency to clarify whether commercially distributed imaging vision-language models, including models marketed with the expectation of later fine-tuning on institutional data, are medical devices requiring premarket review, and to define consistent expectations for validation, performance monitoring and transparency. Notable because a major AI developer is asking for more regulatory scrutiny of its own product category, and because the petition states plainly that downstream users currently bear substantial responsibility for determining whether un-reviewed foundation models are fit for diagnostic use.

Link: https://radiologybusiness.com/topics/artificial-intelligence/radiology-partners-...

Only three of 1,357 FDA-cleared AI devices tested patient outcomes

News-Medical | *20 August 2026*

A PLOS Digital Health analysis of 1,357 FDA-cleared or approved AI and machine learning devices found fewer than 3% linked to a registered clinical trial and only three devices, 0.2%, evaluated against patient outcomes such as morbidity, readmissions or mortality; most public evidence rests on retrospective accuracy or surrogate endpoints. The authors propose a staged validation framework ending in multicenter trials of at least 2,000 participants with patient-centered outcomes and subgroup analysis. DOI 10.1371/journal.pdig.0001597.

Link: https://www.news-medical.net/news/20260820/Only-three-of-1357-FDA-cleared-AI-dev...

SECTION 6

PROFESSIONAL SOCIETIES

SIIM | ACR | RCR | Official publications and event coverage

RSNA R&E Foundation funds 4.5 million dollars across 79 research and education grants for 2026

18 August 2026

The RSNA Research and Education Foundation Board of Trustees approved more than 4.5 million dollars in 2026 grants across 79 projects, including AI-assisted detection of perihilar cholangiocarcinoma, a 5-minute fast-MRI breast protocol, whole-body MRI to study GLP-1 effects, and the foundation's first funded project in Argentina, an open-source AI-assisted simulator for radiology training. Applications for most 2027 grants open in October.

Source: https://www.rsna.org/news/2026/august/re-foundation-fuels-research

ACR prepares a radiology response to the FDA generative AI discussion paper and seeks member input

19 August 2026

One day after the FDA published its discussion paper on regulating generative AI-enabled medical devices, the ACR announced it is preparing a coordinated radiology response through the Data Science Institute and invited members to contribute feedback for possible inclusion in the College's comments before the 19 October deadline. This is the channel through which radiology's operational concerns, validation burden, monitoring obligations and human oversight design, reach the docket.

Source: https://www.acr.org/News-and-Publications/2026/input-on-GenAI

SECTION 7

KEY OPINION LEADERS

LinkedIn | Public statements | Week of 16 to 22 August 2026

Woojin Kim

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

Authored post (16 August) on new behavioral evidence that AI advice suppresses people's willingness to say 'I don't know': across five experiments with 3,132 participants (Marcoccia, Quattrocciocchi, Capraro), the rate of declining to answer dropped from 44% to 3% once AI advice was available, even though the advice was engineered to be wrong, and participants' confidence nearly doubled while accuracy barely moved. Kim's framing for radiology: AI suggestions may alter the metacognitive threshold at which clinicians decide they know enough to answer, which is precisely the automation-bias mechanism reader studies keep surfacing.

Source: LinkedIn

Curt Langlotz

Professor of Radiology, Medicine, and Biomedical Data Science, Stanford; Director, Stanford AIMI

Curation week (18 and 19 August): amplified the FDA Digital Health Center of Excellence's release of its discussion paper on regulating generative AI-enabled medical devices, reposting director Rick Abramson's announcement inviting comment, and relayed the call for papers for AIM-NeurIPS 2026, the first NeurIPS workshop on agentic intelligence for medical imaging and multimodal clinical data (Atlanta, 12 December 2026, submissions due 29 August).

Source: LinkedIn

Amine Korchi

Neuroradiologist, Geneva; HealthTech innovation and ventures

Authored post (21 August) announcing a new paper with Jan Beger and Christoph Agten in European Journal of Radiology Artificial Intelligence: 'Three futures for the diagnostic radiologist: a structured disagreement about what AI actually changes', laying out three scenarios for the specialty, expansion, concentration and stratification, as a structured way to disagree about what AI changes in radiology rather than trading slogans about replacement.

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 16 to 22 August 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.