DR. SERGEY MOROZOV

WEEKLY AI INTELLIGENCE REPORT

Radiology and Medical Imaging

Week 29 | 12 to 18 July 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	6	4	2
Peer-reviewed papers	Industry / FDA items	Media highlights	Key opinion leaders

SECTION 0

EXECUTIVE SUMMARY

Key trends, week of 12 to 18 July 2026

[THEME] RENAL CANCER DOMINATES THE WEEK WITH FOUR MULTIMODAL FUSION MODELS

Four of the eight selected papers target renal disease, all sharing one recipe: fuse imaging with pathology or clinical data and interpret with SHAP. Zhang Y et al. (Academic Radiology) reach AUC 0.948 internally for early renal fibrosis on multimodal ultrasound, dropping to 0.823 externally. Chen X et al. (npj Precision Oncology) subtype RCC on CT with a YOLOv11 system at 0.24 s per file. Chen S et al. (npj Digital Medicine) read a CD8+ T cell signature straight from H&E slides (AUC 0.781). Wang C et al. (J Imaging Inform Med) fuse CT, biopsy biomarkers and clinical data for ccRCC diagnosis and survival.

[EMERGENCY] AI QUANTIFIES FREE AIR ON CT WITH NEAR-PERFECT VOLUME AGREEMENT

Hu Y et al. (Insights into Imaging) report a deep-learning model for pneumoperitoneum on abdominal CT across 2072 patients: AUC 0.97 in testing and sensitivity 84.3% with specificity 89.6% in a 214-scan external emergency cohort. Sensitivity rises to 96% once free-air volumes below 1 mL are excluded, and AI-derived volumes match the reference standard almost exactly (ICC 0.996). The novel contribution is a reproducible volume, which turns a binary alert into a potential triage threshold, though trace-volume detection remains the honest weak point.

[PEDIATRICS] TRANSFORMER STRATIFIES REFRACTORY PNEUMONIA FROM EXISTING SCANS

Ying Z et al. (BMC Medical Imaging) build a transformer framework for refractory Mycoplasma pneumoniae pneumonia in 1224 children, holding AUC near 0.89 across two external cohorts and 0.87 in outpatient settings. It uses chest CT that was already clinically indicated rather than mandating new imaging, keeping it inside existing pathways. Notably it does not beat a multimodal nomogram, so its case rests on simplicity and speed, not a performance ceiling.

[EVIDENCE] NO FLAGSHIP FOUNDATION-MODEL PAPER; DESIGNS ARE RETROSPECTIVE AND CHINA-BASED

This week's pool is unusually homogeneous: single and multi-center Chinese cohorts, retrospective designs, and the recurring CT-plus-pathology-plus-SHAP template. External AUCs generally hold in the high 0.8s, but cohort sizes are sometimes absent from abstracts (Wang C, Wang J), and several report the realistic internal-to-external performance drop. The standout for prospective rigour is a protocol, not a result: An B et al. (JMIR Research Protocols) write an AI malignancy risk index into an RCT endpoint panel reading out in 2028.

[WATCHPOINT] EXPLAINABILITY IS NOW TABLE STAKES, NOT A DIFFERENTIATOR

Seven of the eight peer-reviewed selections this week are built around an interpretability layer (SHAP or Grad-CAM), signalling that explainability has shifted from a selling point to a baseline expectation in radiology-AI publishing. The open question is whether attribution maps that align with known pathophysiology, as several papers show, actually change clinician trust and uptake, or whether they mainly satisfy reviewers. On present evidence the field is converging on a shared multimodal-fusion-plus-explainability template, which raises reproducibility but narrows methodological diversity.

SECTION 1

PEER-REVIEWED PAPERS

Source: PubMed | Verified indexing | No preprints | 12 to 18 July 2026

Explainable Machine-learning Model Based on Multimodal Ultrasound for Non-invasive Detection of Early Renal Fibrosis: A Multicenter Study.

Zhang Y, et al. | ***Academic radiology*** | *2026-07-14*

Interpretable multimodal ultrasound machine-learning model for non-invasive detection of early renal fibrosis in chronic kidney disease, combining conventional ultrasound, ultra micro angiography, shear-wave elastography, super-resolution ultrasound and clinical variables. Prospective, multicenter study: 369 participants (161 healthy controls, 208 mild fibrosis) recruited from 16 institutions, with mild fibrosis defined histopathologically. Six ML algorithms were compared, with held-out centers for external validation and SHAP for interpretability.

Key metrics: Internal dataset: the LightGBM fusion model integrating super-resolution ultrasound, shear-wave elastography, conventional ultrasound and clinical variables reached AUC 0.948 (accuracy 0.880, sensitivity 0.987). External held-out centers: AUC 0.823 (accuracy 0.748, sensitivity 0.883). The fusion model beat the clinical baseline in both settings (internal AUC 0.732, external AUC 0.650). SHAP flagged vessel density and fractal dimension alongside serum creatinine and eGFR as top predictors.

Clinical relevance: Early renal fibrosis is normally a biopsy diagnosis, so a fully non-invasive ultrasound surrogate is attractive if it holds up. The prospective multicenter design and the near-19-point external AUC drop (0.948 to 0.823) are both worth flagging: the drop is the realistic number a department would see, and it is large enough that this reads as risk stratification support rather than a standalone test. SHAP aligning imaging features with known fibrosis pathophysiology strengthens the case that the model is not just fitting site artefacts.

PMID 42448485 **DOI** 10.1016/j.acra.2026.06.037

An explainable multimodal 2.5D deep learning-radiomics model for predicting extranodal extension in lung adenocarcinoma using preoperative CT: a multicenter retrospective cohort study.

Wang J, et al. | ***BMC medical imaging*** | *2026-07-15*

Explainable multimodal 2.5D deep-learning radiomics model (RAIEm) predicting extranodal extension in lung adenocarcinoma from preoperative CT, to support risk stratification and treatment planning before surgery. Tumors were auto-segmented with a 3D U-Net using virtual adversarial training; radiomic features came from the intratumoral region and concentric peritumoral rings at 2, 4, 6 and 8 mm, and 2.5D deep-learning features from five axial slices centered on the largest tumor section. Multicenter retrospective cohort with separate training and external validation cohorts, interpreted with SHAP.

Key metrics: In the validation cohort, AUC was 0.78 (95% CI 0.66 to 0.89), with decision-curve analysis showing greater net benefit across most threshold probabilities. RAIEm significantly stratified recurrence-free survival in both training and external validation cohorts (both $p < 0.01$), and ENE-positive patients had significantly shorter recurrence-free survival. SHAP ranked deep-learning features (one 2D and two 2.5D) among the top three contributors. No total cohort size is stated in the abstract.

Clinical relevance: Extranodal extension is a hard preoperative call that meaningfully changes surgical and adjuvant decisions, so even a moderate AUC of 0.78 can be clinically useful if it holds prospectively. The peritumoral-ring design is a sensible attempt to capture the invasive margin biology that pure intratumoral features miss. The caution is the wide confidence interval and the missing cohort size in the abstract; the authors are explicit that prospective multicenter validation is required before routine use.

PMID 42458284 **DOI** 10.1186/s12880-026-02584-w

A Multimodal, Multitask Prediction Framework for Diagnosis and Prognosis of Clear Cell Renal Cell Carcinoma.

Wang C, et al. | ***Journal of imaging informatics in medicine*** | *2026-07-13*

Multimodal, multitask framework for clear cell renal cell carcinoma that jointly predicts diagnosis and prognosis by fusing preoperative CT radiomics, pathology-derived AMACR/P504S biomarker data from preoperative biopsy, and clinical variables. The model was restricted to pathologically confirmed ccRCC (excluding papillary and chromophobe subtypes). Multicenter retrospective study with patients assigned by hospital site into a training cohort and an independent external test cohort, using post-feature fusion to output tumor classification and postoperative survival simultaneously, with SHAP for feature attribution.

Key metrics: External validation: diagnostic discrimination for P504S/AMACR reached AUC 0.983, and prognostic performance for postoperative outcomes reached a C-index of 0.804. The abstract reports supporting calibration curves, decision curve analysis and bootstrap confidence intervals, and subgroup analyses by grade and stage; SHAP indicated complementary contributions from imaging, pathology and clinical inputs. No numeric cohort size is given in the abstract.

Clinical relevance: This is the fourth renal paper in the same week and shares the field's dominant recipe: fuse CT radiomics with pathology and clinical data, then interpret with SHAP. The external AUC of 0.983 for the biomarker task is high enough to warrant caution about optimism, particularly since the abstract omits cohort sizes, which limits how far the number can be trusted. The multitask design (diagnosis plus survival from one model) is the genuinely useful idea; the authors themselves call for larger prospective multicenter validation before any clinical claim.

PMID 42443645 **DOI** 10.1007/s10278-026-02108-6

Deep learning-based CD8+ T cell model for predicting prognosis and targeted immunotherapy benefit in ccRCC.

Chen S, et al. | ***NPJ digital medicine*** | *2026-07-13*

Deep learning CD8+ T cell inflammation signature (DL-CD8T) for clear cell renal cell carcinoma derived directly from routine H&E-stained whole-slide images, using a clustering-constrained attention multiple-instance learning method to read the tumor microenvironment from standard pathology. Multi-center cohort study with a training cohort and independent external validation in the CPTAC cohort, linking the imaging signature to immune infiltration, checkpoint expression and treatment response.

Key metrics: DL-CD8T distinguished higher CD8+ T cell inflammation from H&E slides with AUC 0.781 (0.741 to 0.817, $p < 0.001$) in training and 0.741 (0.599 to 0.853, $p < 0.001$) in the independent CPTAC cohort. DL-CD8T-positive status was associated with higher CD8+ T cell infiltration; higher PD-1, PD-2, PD-L1, CTLA-4 and TIM-3 expression; and higher tumor mutation burden. It predicted combined targeted and immunotherapy benefit with a hazard ratio of 0.27 ($p = 0.029$).

Clinical relevance: The appeal is economic as much as clinical: inferring an immune-inflammation phenotype from the H&E slide every ccRCC patient already has avoids the cost of immunohistochemistry or sequencing. AUCs in the high 0.7s are modest, and the abstract is explicit that this is pending prospective validation, so the honest framing is a low-cost screening layer that flags who might benefit from combined therapy rather than a decisive predictive biomarker. The HR of 0.27 for treatment benefit is the most interesting signal and the one most in need of confirmation.

PMID 42443348 **DOI** 10.1038/s41746-026-02955-1

Clinical Study on the Prevention of High-Risk Pulmonary Nodule Progression With Yifei Sanjie Pill: Protocol for a Multicenter Randomized Controlled Trial.

An B, et al. | ***JMIR research protocols*** | *2026-07-15*

Protocol for a multicenter randomized controlled trial testing whether Yifei Sanjie Pill, an 8-herb traditional Chinese medicine formula, added to guideline-based CT surveillance reduces progression of high-risk subsolid pulmonary nodules. AI enters as one of several secondary endpoints: an artificial-intelligence-based malignancy risk index tracked alongside standard imaging follow-up. Adults aged 18 to 80 with mixed ground-glass nodules of 8 mm or less are randomized 1:1 to formula granules or matched placebo, twice daily for 6 months, with both arms receiving guideline CT surveillance.

Key metrics: Primary outcome is the 2-year nodule progression rate, defined as longest-diameter increase of 2 mm or more, solid-component increase of 1 mm or more, or a new solid component. As of April 2026, 589 participants had completed randomization; recruitment began June 2023, final 2-year follow-up is scheduled for December 2027, and database lock and primary analysis are planned for early 2028. No performance metrics exist yet because this is a protocol; the AI malignancy risk index is a secondary measure, not a validated result.

Clinical relevance: This is a registered protocol, not a results paper, and it is included as a design signal rather than evidence: AI-based malignancy scoring is now being written into RCT endpoint panels as a standard longitudinal measure of nodule behaviour. That normalisation matters for how imaging AI accrues prospective evidence, since most current radiology-AI validation is retrospective. The honest caveat is that no efficacy or AI-performance number will exist until at least 2028, so the entry weighs future methodology, not present outcomes.

PMID 42456132 **DOI** 10.2196/78534

Deep learning-based segmentation and subtype prediction of renal cell carcinoma on contrast-enhanced CT.

Chen X, et al. | ***NPJ precision oncology*** | *2026-07-14*

Automated two-stage deep learning system that both segments the tumor and predicts renal cell carcinoma subtype on contrast-enhanced CT, built on a YOLOv11 framework for preoperative subtyping. The segmentation model trained on 245 scans from Nanfang Hospital plus 210 from the public KiTS19 dataset; the classification model was developed and validated on 750 patients total, an internal Nanfang cohort (553; 328 training, 112 validation, 113 testing) and two external cohorts (Beijing Tongren, n = 111; two further centers combined, n = 86).

Key metrics: Discrimination of clear cell RCC: AUC 0.878 (internal validation), 0.892 (internal testing), 0.911 (external set I) and 0.892 (external set II). Computational efficiency 0.24 s per file, with FLOPs reduced four times versus conventional 3D CNNs. The abstract notes future work should target sensitivity for small lesions.

Clinical relevance: The distinctive claim here is efficiency, not just accuracy: a quarter-second per file and a fourfold FLOPs reduction is an argument about deployability on ordinary hardware, which is often the real barrier to CT AI adoption. External AUCs around 0.89 to 0.91 across three separate sites are consistent enough to take seriously. The stated blind spot, sensitivity for small lesions, is exactly where preoperative subtyping is hardest, so this is best read as a triage aid rather than a replacement for histology.

PMID 42448906 **DOI** 10.1038/s41698-026-01609-5

Diagnostic performance of an artificial intelligence algorithm for detecting pneumoperitoneum on abdominal CT scans.

Hu Y, et al. | ***Insights into imaging*** | *2026-07-18*

Deep learning model for automated detection, segmentation and volumetric quantification of pneumoperitoneum (free intraperitoneal air) on abdominal CT, aimed at speeding emergency triage where small-volume free air is easy to miss. Multi-center development on 2072 patients split roughly 7:3 into training and testing, with a separate external validation set of 214 emergency CT scans acquired between April 2022 and December 2024. Diagnostic reports served as the reference standard.

Key metrics: Test set (n = 607): sensitivity 91.4%, specificity 93.1%, AUC 0.97 (95% CI 0.95 to 0.99). External emergency cohort (n = 214): sensitivity 84.3% (95% CI 76.2 to 90.5), specificity 89.6% (95% CI 82.3 to 94.6), accuracy 86.9% (81.6 to 91.2), PPV 89.2%, NPV 84.8%. After excluding cases with minimal free gas below 1 mL, sensitivity rose to 96%. AI-derived free-air volumes agreed closely with reference volumes (ICC 0.996, 95% CI 0.994 to 0.997).

Clinical relevance: Pneumoperitoneum is a time-critical finding where a missed or delayed read carries direct surgical consequences, so a model that both flags and quantifies free air maps cleanly onto an emergency worklist-prioritisation use case. The honest weak point is stated in the abstract: trace-volume detection remains hard, and sensitivity only reaches 96% once sub-1 mL cases are set aside. The near-perfect volume ICC is the more novel contribution, because a reproducible volume opens the door to triage thresholds rather than a binary alert.

PMID 42471462 **DOI** 10.1186/s13244-026-02348-8

A CT-based deep learning model for the automated risk stratification of refractory Mycoplasma pneumoniae pneumonia in children.

Ying Z, et al. | ***BMC medical imaging*** | *2026-07-17*

Transformer-based deep learning framework (trans-DLF) that risk-stratifies refractory Mycoplasma pneumoniae pneumonia (RMPP) in children from clinically indicated non-contrast chest CT, at the decision point where identifying refractory cases early changes management. Multicenter retrospective cohort of 1224 pediatric patients: training (n = 506), validation (n = 140), internal testing (n = 139), plus two independent external cohorts (n = 331 and n = 108). Benchmarked against a 3D-CNN, a clinical model and a multimodal nomogram, with Grad-CAM for interpretability.

Key metrics: Median age 6.83 years (IQR 5.0 to 8.6); 609 of 1224 (49.8%) male. AUC 0.97 (95% CI 0.96 to 0.98) training, 0.91 (0.86 to 0.96) validation, 0.90 (0.85 to 0.95) internal testing, and 0.89 (0.84 to 0.94) and 0.89 (0.82 to 0.95) in the two external cohorts. It significantly outperformed the clinical model ($p < 0.001$), while its AUCs were not significantly different from the multimodal nomogram. Outpatient-setting AUC 0.87. Grad-CAM localised predictions to consolidations.

Clinical relevance: The design choice that matters is using CT that has already been ordered rather than mandating a new scan, which keeps the tool inside existing paediatric pathways instead of adding radiation. External AUCs holding near 0.89 across two separate sites is a reasonable generalisation signal for a single-modality model. The candid limitation is that it does not beat the multimodal nomogram, so the case for the imaging-only model rests on simplicity and speed at the point of care, not on a performance ceiling.

PMID 42469676 **DOI** 10.1186/s12880-026-02590-y

SECTION 2

INDUSTRY & REGULATION

Sources: Official registers | Press releases | FDA databases

[FDA] Provect AI Secures FDA 510(k) Clearance for Software Converting C-Arm X-Rays into Intraoperative 3D Imaging

Provect AI | 16 July 2026

Provect AI received FDA 510(k) clearance for its 3D-Anywhere software, which reconstructs volumetric 3D images from standard C-arm X-ray systems for spine and orthopedic procedures without a CT scanner or new hardware, exporting to PACS via DICOM. The company also raised 7 million dollars led by Arteria Capital and ValueStream Ventures.

Source: Provect AI

[REGULATION] Topcon Healthcare Receives CE Mark for GAIA Ultra-Widefield Retinal Imaging Device

Topcon Healthcare | 17 July 2026

Topcon Healthcare announced CE mark approval for GAIA, a fully automated ultra-widefield retinal imaging device offering a 209-degree field of view in a single capture with automated alignment and non-mydriatic operation. Commercial availability in the EU will follow over the coming weeks, with a showcase planned at ESCRS in London in September 2026.

Source: Medical Device News Magazine

[MARKET] Raidium Launches AI-Native Oncology Imaging Viewer R.Read in the US, Deployed at Moffitt Cancer Center

Raidium | 16 July 2026

Raidium announced the US launch of Raidium Read (R.Read), an AI-native oncology imaging viewer offering whole-body lesion detection, AI segmentation and longitudinal lesion transfer across prior studies, now deployed at Moffitt Cancer Center where it replaced a legacy radiomics tool. It is available for clinical trial and research use, with 510(k) clearance for a feature subset expected before year-end 2026.

Source: PR Newswire

[PARTNERSHIP] Nucs AI and Segmed Partner to Build a Validation Platform for Imaging Biomarkers in Oncology and Theranostics

Nucs AI | 14 July 2026

Nucs AI and Segmed announced a strategic partnership, including a Segmed investment in Nucs AI, giving Nucs AI access to de-identified imaging and clinical data spanning more than 2,800 healthcare sites and 150 million imaging studies to develop and validate predictive imaging biomarkers across molecular imaging, radiopharmaceuticals and theranostics.

Source: PR Newswire

[PARTNERSHIP] GE HealthCare and Catholic Health Announce 500 Million Dollar Imaging Care Alliance

GE HealthCare | 16 July 2026

GE HealthCare and Catholic Health (Long Island, New York) announced a 10-year strategic partnership worth about 500 million dollars covering equipment, service, digital products and AI-enabled technologies across more than 40 sites, including MR, CT and PET systems with on-device AI for oncology and cloud-based radiology operations tools such as Imaging 360.

Source: AuntMinnie

[MARKET] Deepnoid Commercializes Generative Medical AI M4CXR After Regulatory Approval

Deepnoid | 13 July 2026

Deepnoid announced the commercial launch of M4CXR, a generative AI chest X-ray device that produces a preliminary findings report covering more than 41 types of findings, after obtaining digital medical device product approval in South Korea. The company said it plans to expand to chest CT and MRI and move toward a medical AI agent service.

Source: The Herald Business

SECTION 3

MEDIA HIGHLIGHTS

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

Radiologists Can Identify Only Three Out of Four AI-Generated Images

AuntMinnie | *17 July 2026*

Covering a Clinical Radiology study by Cronshaw and Williams (University of Edinburgh), AuntMinnie reports that 182 radiologists correctly identified AI-generated images 75.0 percent of the time versus 83.4 percent for real images, with CT and MR the hardest to distinguish. The authors warn that synthetic images could aid training but also risk misuse, fake images in the literature, and misdiagnosis if they enter patient records.

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

The Downside of AI: Some Radiologists Accept Recommendations Even When They Are Incorrect

Radiology Business | *15 July 2026*

Reporting a study in RSNA's journal Radiology, Radiology Business describes automation bias in screening mammography: across 10 breast radiologists read twice six weeks apart with eye-tracking, median sensitivity on false-negative AI suggestions fell to 39 percent with AI versus 71 percent without. An accompanying editorial notes less experienced readers are more prone to over-reliance.

Link: https://radiologybusiness.com/topics/artificial-intelligence/downside-ai-some-ra...

Can AI Solve Radiology's Efficiency Crisis?

ITN Online | *16 July 2026*

In an ITN feature, Curtis Langlotz (Stanford AIMI) argues detection and triage tools can slow radiology by generating findings to adjudicate, while reporting tools give radiologists a draft rather than a blank page; Rad AI co-founder Jeff Chang cites roughly one hour saved per nine-hour shift. Langlotz frames reporting and agentic tools as workflow enablers with humans kept in the loop.

Link: https://www.itnonline.com/article/can-ai-solve-radiology%E2%80%99s-efficiency-cr...

SCCT: AI Detection of Coronary Plaque on Chest CT Increases Lipid-Lowering Therapy

Diagnostic Imaging | *16 July 2026*

Diagnostic Imaging reports Brigham and Women's results from the multicenter AI INFORM study (Nanox AI), where notifying physicians of AI-detected coronary plaque on chest CT drove roughly 33 percent and 46 percent increases in lipid-lowering therapy at 6 and 12 months, with a mean LDL-C drop of 17 mg/dL versus 4 mg/dL in usual care. A separate 5,468-scan study showed over 90 percent agreement between the CAC algorithm and radiologists.

Link: https://www.diagnosticimaging.com/view/scct-impact-ai-detection-coronary-plaque-...

SECTION 4

PROFESSIONAL SOCIETIES

SIIM | ACR | RCR | Official publications and event coverage

RSNA News: AI-generated summaries could help patients understand radiology reports

16 July 2026

RSNA News reported a University of California San Francisco study, published in Radiology Advances, in which a secure version of ChatGPT produced simplified summaries of three lung cancer CT screening reports, each checked by a radiologist for accuracy. Across 1,815 survey respondents, the AI summaries improved comprehension and reduced anxiety, and increased willingness to wait for a scheduled follow-up rather than call the same day. The authors urged clinicians to lead rigorous evaluation of these tools.

Source: https://www.rsna.org/news/2026/july/ai-helps-patients-understand-radiology-reports

SECTION 5

KEY OPINION LEADERS

LinkedIn | Public statements | Week of 12 to 18 July 2026

Curt Langlotz

Professor of Radiology, Medicine and Biomedical Data Science, Stanford University School of Medicine

On 15 July reshared Nishith Khandwala announcing that Bunkerhill Health had raised 55 million dollars from Vinod Khosla at Khosla Ventures, Alfred Lin at Sequoia Capital, Y Combinator and other investors. On 12 July reshared Louis Blankemeier, co-founder and CEO of Cognita, who described starting Cognita in late 2024 with Zhihong Chen and Akshay Chaudhari to build a new ambient radiology reporting system intended to replace Nuance PowerScribe and Fluency. Both amplifications spotlight the reporting and ambient-documentation layer as the current commercial battleground in radiology AI.

Source: LinkedIn

Woojin Kim

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

On 12 July reshared an ACR Data Science Institute post asking 'How do you define an AI error?', framing the challenge of evaluating model performance as healthcare AI continues to evolve. The item connects to the ACR push on post-deployment AI quality monitoring, where a shared definition of error is the prerequisite for meaningful concordance measurement between AI outputs and radiology reports.

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 12 to 18 July 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.