DR. SERGEY MOROZOV

# WEEKLY AI INTELLIGENCE REPORT

Radiology and Medical Imaging

**Week 30 | 19 to 25 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 | 4 | 2 | 2 |
|---|---|---|---|
| Peer-reviewed papers | Industry / FDA items | Media highlights | Key opinion leaders |

SECTION 0

## EXECUTIVE SUMMARY

*Key trends, week of 19 to 25 July 2026*

### [RANDOMIZED TRIAL] AI DECISION SUPPORT LIFTS GENETIC ACCURACY IN RETINAL DISEASE

Jia H et al. (Nature Medicine) ran a 300-participant randomized controlled trial of Retina4IRD, an AI decision support system that predicts 17 genotype categories of inherited retinal disease from fundus and OCT images. Top-5 genetic accuracy before sequencing was 88.5% with AI assistance versus 67.3% for specialists alone ($P < 0.001$), and the composite downstream management score was higher (37.7 versus 28.5). It is the rare imaging-AI study with a prospective randomized design and a hard clinical endpoint, not a retrospective AUC.

### [FOUNDATION MODEL] AUDITABLE CHEST X-RAY MODEL BUILDS INTERPRETABILITY INTO ITS REASONING

Han T et al. (Nature Biomedical Engineering) present CLEAR, an auditable foundation model trained on over 0.87 million image-report pairs from 239,391 patients, which projects chest X-rays into a concept space so every prediction decomposes into named radiological observations. Validated across the United States, Europe and Asia, it enables auditable zero-shot detection and confounder identification. The abstract claims state-of-the-art performance without publishing the numbers, so the quantitative edge cannot be judged from it.

### [BIOMARKER] SPEECH MODEL TRACKS HUNTINGTON'S PROGRESSION AND MATCHES CLINICIAN RATINGS

Le Moine Veillon C et al. (Lancet Digital Health) developed NDSNet, a deep learning model that estimates clinical scores from speech in Huntington's disease across three longitudinal cohorts. Relative error was 11.4% and the intraclass correlation with observed motor scores was 0.87, similar to clinician ratings (0.847), with predictions grounded against striatal atrophy on MRI and a signal already present at presymptomatic HD-ISS stages 0-1. External replication was weaker (ICC 0.67), the realistic performance figure.

### [ONCOLOGY] MULTIMODAL FUSION SPREADS ACROSS LIVER, PROSTATE AND ESOPHAGEAL IMAGING

Three oncology papers share the multimodal-fusion recipe. Li F et al. (npj Digital Medicine) built TEMPO-HCC, a time-resolved survival model fusing MRI and pathology (external C-index 0.751). Fu X et al. (Abdominal Radiology) predict positive surgical margins in prostate cancer from MRI plus clinical data (external AUC 0.843 across six centers). Ye C et al. (European Journal of Radiology) detect early esophageal cancer on noncontrast CT (external AUC 0.847), lifting all four readers' sensitivity.

**[WATCHPOINT] INTERPRETABILITY AND AUDITABILITY MOVE FROM FEATURE TO FOUNDATION**

This week interpretability is not a bolt-on. CLEAR (Han T et al., Nature Biomedical Engineering) builds concept-level auditing into the representation itself, TEMPO-HCC (Li F et al., npj Digital Medicine) exposes an auditable macro-to-micro evidence chain, and a Chinese expert consensus (Wu H et al., Insights into Imaging) codifies interpretable outputs and post-deployment surveillance as baseline expectations across 30 statements. Read alongside the trade-press coverage of the EU AI Act, the signal is that auditability is becoming a regulatory and procurement requirement rather than a research nicety, which favours models designed to explain themselves over marginally more accurate black boxes.

SECTION 1

# PEER-REVIEWED PAPERS

*Source: PubMed | Verified indexing | No preprints | 19 to 25 July 2026*

## AI-based clinician decision support system for diagnosis of inherited retinal diseases: a multicenter, randomized trial.

*Jia H, et al.* | ***Nature medicine*** | *2026-07-24*

AI-based clinician decision support system (Retina4IRD) that predicts 17 genotype categories of inherited retinal diseases from retinal images, using a Vision Transformer pretrained with RETFound on color fundus photographs and OCT scans. It was trained and validated on multimodal data from 1,843 genetically confirmed patients (3,376 eyes) across China, South Korea and Poland, then tested in a randomized controlled trial of 300 participants with suspected inherited retinal disease randomized 1:1 to a Retina4IRD-assisted specialist arm or a specialist-only arm, with 295 analysed against next-generation sequencing.

**Key metrics:** Top-5 prediction accuracy 0.904 (95% CI 0.896 to 0.912) internal and 0.856 (0.850 to 0.863) external. In the RCT the primary outcome was met: top-5 genetic accuracy 88.5% in the assisted arm versus 67.3% specialist-only (P < 0.001); top-1 37.8% versus 22.4% and top-4 81.8% versus 53.1%. The composite downstream management score was higher with assistance (37.7 versus 28.5, P < 0.001). Median age 33 years, 114 of 295 (38.6%) female.

**Clinical relevance:** This is the rare imaging-AI study with a prospective randomized design and a hard clinical-accuracy endpoint rather than a retrospective AUC, which is why it tops the week. A 21-point absolute gain in top-5 genetic accuracy before genetic testing maps directly onto the diagnostic odyssey in inherited retinal disease, and the improved downstream management score suggests the benefit extends beyond the label. The honest caveat is that the tool assists rather than replaces specialists and top-1 accuracy stays modest at 37.8%, so it narrows the genetic differential ahead of sequencing rather than delivering a definitive genotype.

**PMID** 42498742 **DOI** 10.1038/s41591-026-04545-w

## CLEAR: an auditable foundation model for radiology grounded in clinical concepts.

*Han T, et al.* | ***Nature biomedical engineering*** | *2026-07-22*

Auditable radiology foundation model, CLEAR (Concept-Level Embeddings for Auditable Radiology), that projects chest X-rays into a semantically rich space defined by large language model embeddings, so every prediction decomposes into weighted contributions from individual radiological observations. It was trained on over 0.87 million image-report pairs from 239,391 patients and externally validated on four large, physician-annotated datasets from the United States, Europe and Asia.

**Key metrics:** The abstract reports state-of-the-art classification performance on external validation but gives no single headline AUC or sensitivity. It instead emphasises capabilities: auditable zero-shot pathology detection, systematic identification of radiological confounders, and the creation of expert-level concept bottleneck models from data-driven concepts. No numeric performance value is stated in the abstract, so the comparative claim cannot be quantified from it.

**Clinical relevance:** CLEAR is the week's flagship methods contribution and it attacks the trust problem at the root: rather than a post-hoc saliency overlay, interpretability is built into the representation so a prediction traces back to named radiological observations. That kind of auditability is exactly what post-market monitoring and the EU AI Act now expect. The caution for readers is that the abstract asserts a state-of-the-art edge without publishing the numbers, so the quantitative advantage over existing chest X-ray models cannot be judged from the abstract alone.

**PMID** 42486902 **DOI** 10.1038/s41551-026-01741-4

## A CNN-based liver ultrasound plane recognition system with multimodal large language model-generated interpretable feedback.

*Chen Y, et al.* | ***Abdominal radiology (New York)*** | *2026-07-22*

Hybrid AI pipeline that pairs a convolutional neural network for image recognition with a multimodal large language model for textual interpretation, to identify liver ultrasound standard planes and generate interpretable feedback. This multicenter study used 16,782 images in the standard dataset (training 13,508; testing 3,274), an internal validation set of 3,431 images and an external validation set of 5,253 images, comparing students, junior, intermediate and senior radiologists, the AI model, and the same readers with AI assistance against expert consensus.

**Key metrics:** The CNN recognised 13 liver ultrasound standard planes and nonstandard planes with 88.8% accuracy in the test set. In validation the model's AUC, sensitivity, specificity, PPV, NPV, accuracy and kappa were similar to senior radiologists (all $p > 0.05$); it outperformed students, junior and intermediate radiologists on all metrics except PPV, and readers improved significantly with hybrid AI assistance (all $p < 0.05$), with robust generalisation across centers and devices.

**Clinical relevance:** Standard-plane acquisition is the operator-dependent bottleneck in ultrasound, so a model that matches senior readers and lifts less-experienced ones is a sensible quality-control and training use case rather than a diagnostic claim. The multimodal LLM feedback layer is the differentiator, turning a plane classifier into a teaching aid. The authors are explicit that this is image quality control, not disease detection, and call for prospective validation before clinical deployment.

**PMID** 42484849 **DOI** 10.1007/s00261-026-05706-8

## Interpretable multimodal deep learning for time-resolved survival prediction after hepatocellular carcinoma resection.

*Li F, et al.* | ***NPJ digital medicine*** | *2026-07-20*

Multimodal deep survival model, TEMPO-HCC, with hierarchical interpretability that estimates individualised overall-survival risk trajectories after curative-intent resection of hepatocellular carcinoma, integrating multiphasic MRI, postoperative H&E whole-slide images and perioperative predictors. It was developed and externally validated on a six-center cohort of 1,475 patients, with a discrete-time survival head generating probabilities at 1, 2, 3 and 5 years after surgery.

**Key metrics:** In external validation TEMPO-HCC reached a C-index of 0.751 and time-dependent AUCs of 0.836, 0.781, 0.812 and 0.680 at 12, 24, 36 and 60 months. It outperformed unimodal models and guideline-based staging systems, and augmenting guideline staging with TEMPO-HCC improved discrimination.

**Clinical relevance:** The novelty is time-resolved risk rather than a static binary label, which matches how post-resection surveillance is actually scheduled. Fusing MRI with histopathology through an auditable macro-to-micro evidence chain is the design several groups are converging on, and the external C-index of 0.751 is a credible, non-inflated number. The weak point is visible in the numbers: the 5-year AUC falls to 0.680, so long-horizon prediction remains the hard part.

**PMID** 42477420 **DOI** 10.1038/s41746-026-03027-0

## A deep learning model for speech-based prediction of clinical scores in people with Huntington's disease: a longitudinal study with cross-sectional replication.

*Le Moine Veillon C, et al.* | ***The Lancet. Digital health*** | *2026-07-23*

Automated deep learning model, NDSNet (Neurodegenerative Disease Speech Network), that generates speech-derived estimates of contemporaneous clinical scores at each visit in Huntington's disease, from presymptomatic (HD-ISS stages 0-1) to symptomatic stages. It combines a pre-trained wav2vec 2.0 model with a recurrent attentive network in a contrastive-learning framework, was developed across three prospective longitudinal cohorts, externally validated in a replication cohort, and compared against striatal atrophy on MRI, the best-established marker of progression.

**Key metrics:** The developmental cohort included 191 people with Huntington's disease and 58 controls; overall relative error between predictions and observed scores was 11.4% (95% CI 9.7 to 12.5). The intraclass correlation between predicted and observed motor scores was 0.87 (0.83 to 0.91), similar to clinician ratings (ICC 0.847). In the replication cohort (67 patients after exclusions) relative error was 15.0% (13.0 to 19.0) and ICC 0.67 (0.38 to 0.71). Predicted scores showed MRI associations similar to observed scores, and stronger associations with striatal atrophy at HD-ISS stages 0-1.

**Clinical relevance:** Speech is an attractive imaging-adjacent biomarker because it is cheap and repeatable, and the neuroanatomical grounding against striatal atrophy is what separates this from a black-box audio classifier. The subclinical-stage signal at HD-ISS 0-1 is the most important claim, since that is where trials need sensitive endpoints before overt brain damage. The honest limitation is the drop in the replication cohort (ICC 0.67 with a wide interval), which is the realistic external-performance figure.

**PMID** 42493284 **DOI** 10.1016/j.landig.2026.101025

## Preoperative prediction of positive surgical margins in prostate cancer using multimodal deep learning model: a multicenter study.

*Fu X, et al.* | ***Abdominal radiology (New York)*** | *2026-07-25*

Multimodal deep learning model predicting the risk of positive surgical margins after radical prostatectomy from preoperative MRI and clinical features. This retrospective multicenter study included 1,177 prostate cancer patients across eight institutions, with 1,022 patients from two institutions for training and 155 from six independent centers for external validation. A feature-disentanglement model (DESM) isolates disease-specific features from hospital-specific variation, and a multimodal fusion model (MDESM) adds clinical variables, with Grad-CAM for interpretability.

**Key metrics:** In external validation MDESM achieved an AUC of 0.843 (95% CI 0.770 to 0.911), significantly outperforming the imaging-only DESM (0.711, 0.607 to 0.800, p = 0.003) and the clinical-only model (0.676, 0.577 to 0.770, p < 0.001). Decision curve analysis showed the highest net benefit for MDESM across a range of threshold probabilities.

**Clinical relevance:** Positive surgical margins drive biochemical recurrence, so a preoperative estimate that could inform nerve-sparing decisions is clinically meaningful. The feature-disentanglement design directly targets the multi-site generalisation failure that sinks most radiomics, and holding an AUC of 0.843 across six external centers is a reasonable signal. The realistic caveat is the retrospective design and a wide external interval (0.770 to 0.911), so surgical-planning support rather than a determinant is the right framing.

**PMID** 42501097 **DOI** 10.1007/s00261-026-05672-1

## Chinese expert concern and consensus on applications of artificial intelligence in clinical cancer imaging.

*Wu H, et al.* | ***Insights into imaging*** | *2026-07-24*

National multidisciplinary expert consensus from China on applications of AI in clinical cancer imaging, developed with a modified Delphi approach. Eighty-one specialists in radiology, nuclear medicine, oncology and imaging AI from university hospitals completed a survey of 30 core statements spanning screening, diagnosis, staging, treatment planning, response assessment, prognosis, data governance and implementation.

**Key metrics:** Consensus was defined as a mean score of at least 7 on a 9-point Likert scale with at least 80% of experts scoring 7 or higher. All 30 statements met the thresholds, with mean scores ranging from 8.06 to 8.58 and the proportion of experts scoring 7 or higher ranging from 86% to 98%.

**Clinical relevance:** This is a governance and implementation document rather than a new model, included as a signal of where consensus is settling: expert-supervised labeling, multicenter validation, subgroup evaluation, interpretable outputs, privacy-preserving collaboration and post-deployment surveillance. Its value is as a checklist for departments and vendors, not as performance evidence. The obvious limitation is that it reflects a single national panel, so its recommendations should be read alongside ACR, ESR and RSNA guidance rather than in place of it.

**PMID** 42496805 **DOI** 10.1186/s13244-026-02359-5

### Integrating habitat-based radiomics, 3D deep learning, and clinical variables for detection of early-stage esophageal squamous cell carcinoma on noncontrast chest CT: A multicenter study.

*Ye C, et al.* | ***European journal of radiology*** | *2026-07-23*

Combined model integrating habitat-based sub-regional radiomics, 3D deep learning features and clinical variables for detecting early-stage esophageal squamous cell carcinoma on noncontrast chest CT. In this retrospective multicenter study the segmented esophageal region was partitioned into habitat sub-regions by unsupervised clustering, and radiomic features were fused with deep features from a 3D ResNet-101 into a deep-learning-radiomics model, then combined with clinical variables, with SHAP for interpretability and an exploratory reader-comparison study.

**Key metrics:** A total of 469 subjects were included. The Combined model achieved the highest AUC of 0.895 in internal validation and 0.847 in the external test, with favorable calibration and decision curve analysis. AI assistance improved the sensitivity of all four radiologists for early-stage disease and increased inter-reader agreement.

**Clinical relevance:** Early esophageal squamous cell carcinoma is nearly invisible on routine noncontrast CT, so an opportunistic second reader on already-acquired scans is a genuinely useful screening angle rather than a new test. The habitat-based sub-regional approach is a sensible attempt to model tumor heterogeneity, and lifting all four readers' sensitivity is the most practically relevant result. The limitation is the modest sample of 469 and the retrospective case-control design, and the authors explicitly call for prospective evaluation.

**PMID** 42501614 **DOI** 10.1016/j.ejrad.2026.113106

SECTION 2

## INDUSTRY & REGULATION

*Sources: Official registers | Press releases | FDA databases*

### [MARKET] CARPL.ai raises USD 10 million Series A led by IFC to scale global radiology AI adoption

*CARPL.ai | 23 July 2026*

The radiology AI marketplace CARPL.ai closed a USD 10 million Series A led by the International Finance Corporation, the World Bank Group's private sector arm, with Stellaris Venture Partners. The platform aggregates more than 300 AI applications from over 100 vendors behind a single integration, letting providers deploy, test and monitor imaging AI across workflows. Funds target product development and expansion across North America, Latin America, Europe and Asia-Pacific.

Source: CARPL.ai

### [FDA] VELMENI receives FDA 510(k) clearance for AI-powered 3D CBCT dental imaging software

*VELMENI | 23 July 2026*

VELMENI obtained FDA 510(k) clearance for VELMENI for DENTISTS (V4D) 3D, extending its dental AI portfolio into cone-beam CT. The software assists dental professionals with AI-based anatomical segmentation, visualization, measurements and reconstruction of 3D CBCT scans, and joins the company's connected 2D and 3D imaging platform.

Source: PR Newswire

### [MARKET] HOPPR launches EF Chest CT Narrative foundation model, its third imaging modality

*HOPPR | 23 July 2026*

HOPPR released its EF Chest CT Narrative Model, a foundation model that ingests 3D chest CT volumes and generates narrative language across pulmonary, mediastinal, cardiac, upper abdominal, osseous and soft tissue regions. It is the company's fifth foundation model and extends its portfolio to a third modality after chest X-ray and mammography, delivered through HOPPR Forward Deployed Services on HIPAA-ready infrastructure.

Source: HOPPR

### [PARTNERSHIP] Telefonica Tech and Harrison.ai partner to deploy CE-marked chest X-ray AI across Spain

*Telefonica Tech | 21 July 2026*

Telefonica Tech partnered with Harrison.ai to roll out an AI chest X-ray solution across Spain's regional health services. The CE class IIb certified algorithm, trained on more than 780,000 chest X-ray studies each labelled by at least three radiologists, flags up to 124 findings including pulmonary nodules for lung cancer screening, integrating into existing workflows to prioritise urgent cases while the clinician retains final say.

Source: Telefonica

SECTION 3

## MEDIA HIGHLIGHTS

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

### LLMs could automate radiology quality control, saving departments significant time

***Radiology Business*** | *23 July 2026*

Reporting on a Journal of the American College of Radiology study, Radiology Business describes how a large language model automated quality control of breast ultrasound reports across 735 patients at 60 hospitals, outperforming manual reviewers on lesion margins (80 percent versus 64 percent) and echo patterns (74 percent versus 56 percent) and completing checks on 50 reports in 13 minutes versus 213 minutes for human reviewers.

Link: https://radiologybusiness.com/topics/artificial-intelligence/llms-could-automate...

### EU AI Act brings radiology AI oversight into daily practice

***AuntMinnieEurope*** | *21 July 2026*

An opinion paper covered by AuntMinnieEurope argues that meeting the EU AI Act's AI literacy (Article 4) and post-market monitoring (Article 72) duties on paper is not enough. It urges radiology departments to build practical competence through documented local audits of deployed models and continuous surveillance for drift and demographic bias, with medical-AI monitoring requirements expected to apply from August 2028.

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

SECTION 4

## PROFESSIONAL SOCIETIES

*SIIM | ACR | RCR | Official publications and event coverage*

SECTION 5

## KEY OPINION LEADERS

*LinkedIn | Public statements | Week of 19 to 25 July 2026*

### Amine Korchi

*Medical Director and Radiologist; HealthTech investor*

Reacting to Sectra's acquisition of Oxipit for just EUR 1.7 million in cash, Korchi argued the low price shows that being first is not enough: Oxipit built the world's first CE-marked autonomous chest X-ray AI, able to clear normal studies from the worklist without human review, but the market was not ready for autonomous AI, and being too early is indistinguishable from being wrong. He still expects autonomous AI to arrive, driven by rising imaging demand and radiologist shortages.

Source: LinkedIn

### Curt Langlotz

*Professor of Radiology and Director, Stanford AIMI Center*

In the opening keynote of RANZCR's Intelligence26 AI conference on 24 July, Langlotz addressed AI's effect on the radiology workforce. As reported from the meeting, he argued that radiology leads medicine in AI adoption but also in predictions of its own extinction, and that AI will shift rather than replace radiologist tasks, absorbing rising imaging volume rather than displacing specialists.

Source: LinkedIn (Croakey Health Media, reporting from RANZCR Intelligence26)

**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 19 to 25 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.*