



Ethical AI in Radiology: Performance, People, and Post-market Responsibility

Author: Alex Perkins, EMJ, London, UK

Citation: EMJ Radiol. 2026;7[1]:23-26.
<https://doi.org/10.33590/emjradiol/19453N4H>



AT THE European Congress of Radiology (ECR) 2026, the session 'The Art of Ethical AI: Redefining Performance in Radiology' brought together speakers who argued that ethical AI in radiology cannot be judged by accuracy alone. Chaired by Elmar Kotter, University of Freiburg, Freiburg im Breisgau, Germany, the session explored how regulation, post-market surveillance, and human factors are reshaping what 'good performance' means in clinical practice. Across the presentations, one message stood out clearly: AI implementation only works when technical performance, governance, and real-world workflow are considered together.

FROM PRINCIPLES TO PRACTICE: WHAT THE EU AI ACT MEANS FOR RADIOLOGY

Hugh Harvey, Hardian Health, London, UK, emphasised that the EU AI Act (2024) is no longer a future policy question, but a framework already shaping how AI is developed and deployed in radiology. Harvey highlighted how the legislation distributes responsibility across the entire AI lifecycle, from development through to real-world use.

A central component is risk management. Under Article 9, high-risk AI systems must have a risk management system that is established, implemented, documented, and maintained. In practice, this requires developers to embed structured processes within their quality management systems, including assessment of performance in real-world environments, as well as safeguards for data protection, cybersecurity, and adverse event reporting.

Transparency is another key pillar. Article 13 requires developers and distributors to provide deployers with clear and comprehensive information, including

instructions for use, human oversight measures, performance characteristics, technical capabilities, and maintenance requirements. For radiology departments, this ensures that AI tools are not treated as 'black boxes,' but as systems whose limitations and operational requirements must be understood in clinical context.

“Hospitals must also take an active role in ensuring safe and appropriate use”

Importantly, Harvey stressed that responsibility does not rest with vendors alone. Article 26 outlines the obligations of deployers, meaning that hospitals must also take an active role in ensuring safe and appropriate use. This includes assigning adequately trained human oversight, monitoring system performance in line with the instructions for use, retaining system logs, and reporting any safety issues to both manufacturers and regulatory authorities.

He also pointed to ongoing efforts to refine how the legislation is implemented

in practice. The proposed 'digital omnibus', introduced in November 2025, represents a series of targeted amendments intended to address criticisms of the act and make its requirements more workable for industry and healthcare providers. Among the proposed changes is a shift in responsibility for AI literacy, removing the direct obligation on providers and deployers to deliver training, and instead placing greater emphasis on guidance and support from the European Commission and Member States.

Alongside this, Article 57 introduces regulatory sandboxes at national level, with the prospect of an EU-wide sandbox expected by 2028. Taken together, these developments suggest that while the regulatory framework is already in place, its practical application is still evolving, with increasing focus on balancing oversight, usability, and innovation in clinical AI.

SURVEILLANCE CANNOT STOP AT DEPLOYMENT

While regulation sets the framework, Kicky Gerhilde Van Leeuwen, Romion Health

& Health AI Register, Utrecht, the Netherlands, emphasised that post-market surveillance is what determines whether AI remains safe once it reaches clinical practice. Her presentation focused on a basic, but often neglected question: how do clinicians ensure long-term safety when AI tools are scaled across dynamic healthcare systems?

Van Leeuwen pointed out that many hospitals still validate AI on their own local datasets before use. In the Netherlands, for example, more than 20 out of 70 hospitals have fracture detection tooling, and each tests performance locally first. While understandable, she questioned whether this approach is sustainable. Clinicians would not routinely retest an approved drug or a new CT scanner in every individual hospital population, so why is AI treated differently?

Part of the answer, she suggested, lies in the fact that AI is uniquely sensitive to change. Imaging hardware changes, post-processing changes, algorithms are updated, and patient populations drift over time. This means evaluation cannot be a one-off event. Instead, she described a





continuum that starts with retrospective analysis of available evidence, moves through acceptance testing and piloting in the local workflow, and then continues into post-deployment monitoring.¹ That need is especially pressing given the evidence base for commercial tools remains uneven.

Van Leeuwen also highlighted a striking regulatory gap: none of the 13 manufacturers visited by the Dutch Health and Youth Care Inspectorate in 2023/2024 met post-market surveillance requirements. In her view, customer surveys are no substitute for structured monitoring. What matters is measurement across several domains: technical metrics, such as uptime and latency; clinical metrics, such as drift and diagnostic performance; and impact metrics, such as user experience and efficiency. As she put it: “If we want to ensure long-term safety of AI, in a world where the only constant is change, we need post-deployment monitoring.”

WHEN AI CHANGES THE SYSTEM, PEOPLE FEEL IT FIRST

Susan Cheng Shelmerdine, University College London, UK, brought the discussion firmly into the human domain, asking

what happens when radiology embraces AI without accounting for workflow, staff experience, and patient expectations. Using the example of AI-supported lung cancer triage, she showed that the technology initially appeared highly successful. Following implementation, the proportion of patients meeting the national target of 72 hours from abnormal chest X-ray to CT increased from 19.2% to 46.5%, same-day CT rose from 4.0% to 22.1%, and the average time from X-ray to CT fell from 6 days to 3.6 days.²

Yet, this apparent success did not guarantee sustainability. Shelmerdine explained that the pathway was later withdrawn in 2025, illustrating a central lesson of implementation science: adding AI changes the entire system, not just one reporting step. Drawing on the Systems Engineering Initiative for Patient Safety (SEIPS) model, she showed how AI can disrupt relationships between tasks, people, technologies, and organisational processes.³ Early staff feedback reflected this tension. Before deployment, there was cautious optimism, mixed with uncertainty around deskilling, job security, and patient benefit. One month after deployment, feedback turned largely negative, centring on workflow disruption, delays, false

positives, and frustrated patients. By 8 months, views had become more balanced, with staff recognising benefits for same-day CT and patient care, while still calling for better communication and pathway design.⁴

“**Ethical AI in radiology is less about whether a tool works in principle and more about whether it can continue to work safely, transparently, and acceptably in practice**”

This evolution, Shelmerdine argued, reflects the gap between ‘work as imagined’ and ‘work as done.’⁵ That gap should not automatically be viewed as failure, but as a signal that systems behave differently under real-world constraints. She linked this to what she has previously described as the cycle of over-investment, honeymoon, disinvestment, and eventual reinvestment that often characterises healthcare’s relationship with AI.⁶

Her presentation also addressed the problem of trust. Clinicians may over-trust, under-trust, or appropriately calibrate their reliance on automation, and those patterns can directly affect

performance.⁷ She cited recent evidence suggesting that even when experts are shown detailed feedback about their own performance and an AI system’s strengths and weaknesses, this does not necessarily translate into substantially improved use of AI support.⁸ For Shelmerdine, the lesson was not that AI should be abandoned, but that responsible implementation must also include responsible withdrawal: clear communication with stakeholders, maintenance of human skills during deployment, and an exit strategy before adoption begins.

CONCLUSION

Taken together, the session suggested that ethical AI in radiology is less about whether a tool works in principle and more about whether it can continue to work safely, transparently, and acceptably in practice. Regulation may define obligations, but it is post-market surveillance and attention to human factors that determine whether those obligations translate into better care. As radiology moves further into AI-enabled practice, success could increasingly depend not on adopting more systems, but on building systems that clinicians, patients, and regulators can realistically live with.

References

1. Antonissen N et al. Artificial intelligence in radiology: 173 commercially available products and their scientific evidence. *Eur Radiol.* 2026;36(1):526-36.
2. Storey M et al. Early clinical evaluation of AI triage of chest radiographs: time to diagnosis for suspected cancer and number of urgent CT referrals. *NEJM AI.* 2025;3(1):DOI:10.1056/Alcs2500539.
3. Carayon P et al. Work system design for patient safety: the SEIPS model. *Qual Saf Health Care.* 2006;15(Suppl 1):DOI:10.1136/qshc.2005.015842.
4. Togher D et al. Evolution of radiology staff perspectives during artificial intelligence (AI) implementation for expedited lung cancer triage. *Clin Radiol.* 2025;DOI:10.1016/j.crad.2024.09.010.
5. Hollnagel E, *Safety-I and Safety-II: The Past and Future of Safety Management* (2014) 1st edition, Oxfordshire: Routledge.
6. Shelmerdine SC. Rethinking our relationship with AI: for better or worse, richer or poorer? *Eur Rad.* 2024;DOI:10.1007/s00330-024-11007-9.
7. Parasuraman R, Manzey DH. Complacency and bias in human use of automation: an attentional integration. *Hum Factors.* 2010;52(3):381-410.
8. Chen C et al. Can domain experts rely on AI appropriately? A case study on AI-assisted prostate cancer mri diagnosis [Internet] (2025) Ithaca: arXiv. Available at: <https://arxiv.org/abs/2502.03482>. Last accessed: 26 March 2026.