Epistemic Authority and Clinical Governance in AI-Integrated Radiography: A Practice-Oriented Narrative Review

Giuseppe Scappatura[1]

1. *Radiology Department, G.O.M. "Bianchi-Melacrino-Morelli", Reggio Calabria, Italy*

Correspondence: giuseppe.scappatura@ospedalerc.it

KeyWords:

Artificial intelligence; radiography; epistemic authority; clinical governance; radiomics; Critical AI Literacy; EU AI Act.

Abstract

Artificial intelligence in radiography reconfigures professional accountability, radiomic data stewardship, and clinical governance obligations across the imaging workflow. This narrative review proposes epistemic authority — the radiographer's accountable role as guarantor of the validity, interpretability, and governance of AI-mediated outputs — as the unifying framework for understanding these transformations. AI integration redistributes procedural tasks without transferring professional responsibility: radiographers retain non-delegable accountability for radiation justification, technical–clinical compromise, and patient-centred judgement, particularly in paediatric, frail, trauma, and non-standard contexts where algorithmic quality thresholds conflict with ethical proportionality. Radiomic validity is contingent on acquisition reproducibility; variability in kVp, mAs, reconstruction algorithm, and voxel dimensions generates epistemically unstable feature distributions, positioning technical standardisation as a form of scientific accountability rather than mere protocol compliance. Systemic vulnerabilities — automation bias, professional deskilling, dataset inequity, and algorithmic drift — are compounding, longitudinal threats requiring governance architectures designed around technical–clinical–organisational interfaces. The EU AI Act formally classifies healthcare AI as high-risk; sustainable compliance requires active professional operationalisation, not passive institutional adoption. Epistemic authority is the defining competency of contemporary radiographic practice, and its cultivation — through Critical AI Literacy embedded in pre-registration curricula, continuing professional development, and institutional governance — represents a genuine expansion of professional authority, not a diminution of it.

Introduction

Beyond Operational Familiarity — The Epistemic Stakes of AI in Radiography

Artificial intelligence in radiography has attracted substantial scholarly attention, with recent contributions addressing professional identity [5, 6], governance frameworks [21], and the perceived impact of automation on radiographers' careers [27]. What this literature has not yet synthesised is the unifying concept that connects these concerns: epistemic authority — the radiographer's accountable role as guarantor of the validity, interpretability, and governance of the AI-mediated outputs that increasingly co-produce diagnostic knowledge.

This distinction matters practically. A radiographer who understands that AI improves throughput, but does not understand that their acquisition decisions determine the scientific validity of radiomic predictions, is operationally trained but epistemically unprepared. A profession that acknowledges governance as important, but treats it as an institutional function rather than an individual competency, will produce practitioners who comply with policy but cannot identify algorithmic drift, challenge boundary failures, or meaningfully participate in procurement and post-market surveillance. The gap in the current literature is not descriptive — the landscape of AI in radiography has been mapped — but integrative: no existing review has articulated the framework that unifies professional role transformation, data stewardship, and governance as expressions of a single, expanded professional accountability.

Artificial intelligence cannot be reduced to a technological upgrade. It represents a structural reconfiguration of how diagnostic knowledge is generated, validated, and operationalised within clinical imaging environments [1, 10, 27]. Conventional radiographic technologies enhanced acquisition through deterministic improvements in detector sensitivity, dose modulation, and digital processing. AI introduces probabilistic inference — algorithmic reasoning derived from large-scale data modelling embedded directly within routine clinical decision pathways — creating what must be understood as a human–algorithmic system in which professional judgement and computational inference co-produce diagnostic artefacts [1, 10, 16, 21]. The radiographer remains the final guarantor of radiation justification, clinical adequacy, and patient safety even where algorithms provide positioning guidance, quality flags, or triage prioritisation [1, 8, 16, 17, 21, 25]. The central question is not whether AI improves measurable performance, but how algorithmic mediation reshapes accountability, expertise, and the boundaries between delegable and non-delegable professional judgement.

Many current AI implementations operate as invisible computational layers: automated centring modules, artefact detection systems, protocol optimisation

JAHC | JOURNAL OF ADVANCED HEALTH CARE
OFFICIAL JOURNAL OF TSRM PSTRP
RESEARCH ARTICLE

engines, clinical triage algorithms [1, 10, 21, 27]. Their invisibility renders them consequential; operating without continuous interrogation, these systems risk becoming epistemically naturalised — accepted uncritically, their influence on clinical perception ungoverned. Epistemic authority, as developed throughout this review, is precisely the professional capacity to resist that naturalisation: to interrogate AI outputs, identify boundary conditions, maintain acquisition standards that preserve data validity, and participate actively in the governance structures that make clinical AI safe across its deployment lifecycle.

Narrative Review Approach

This narrative review was developed through structured searches of PubMed/MEDLINE, Scopus, and CINAHL (2015–2026), complemented by key policy, regulatory, and consensus documents [1, 10, 17, 21, 25, 27]. Search terms included: Artificial Intelligence, Radiographer, Deep Learning, Clinical Governance, Radiomics, Explainability, Automation Bias, and Deskilling. Inclusion criteria required relevance to radiographer practice, imaging workflows, professional responsibility, AI governance, or radiomic reproducibility. The synthesis is organised around five practice-relevant domains: professional role transformation and ethical judgement; radiomic stewardship and acquisition standardisation; performance metrics and systemic vulnerabilities; governance, Critical AI Literacy, and regulatory obligations; and educational and workforce implications [1, 11, 16, 21, 27].

From Technical Operator to Supervisory Clinician

Automation redistributes procedural labour but does not transfer professional responsibility [1, 16, 21]. As radiographic systems assume greater autonomy over positioning, quality assurance, post-processing, and triage, the oversight function of the radiographer becomes more consequential, not less — concentrated around moments of uncertainty, ethical compromise, and patient-specific adaptation that fall outside the competence of probabilistic models [1, 8, 11, 16].

Ethical Judgement in Technical–Clinical Trade-offs

AI quality assurance systems evaluate images against learned statistical norms, not ethical proportionality or patient burden [1, 10, 27]. In paediatric, geriatric, trauma, or frail populations, repeating an exposure to satisfy an algorithmic threshold may be clinically unjustified: a technically optimal image is not necessarily the ethically preferable one when repetition increases radiation burden, pain, or anxiety without meaningful diagnostic gain [1, 8, 16]. Radiation justification and the technical–clinical compromise therefore remain irreducibly human responsibilities — the precise domain where professional authority is most resistant to algorithmic substitution, and where statistical criteria must yield to ethical proportionality [1, 16, 17, 21].

Contextual Interpretation and the Triage Paradox

Machine learning models are constrained by the representativeness of their training data [5, 11–15]. Out-of-distribution cases — atypical anatomy, post-surgical alteration, metallic implants, or acquisition artefacts — may generate confident but clinically misleading outputs.

Case illustration: An automated pneumothorax triage system flags a linear density as a possible pleural line. A radiographer with Critical AI Literacy identifies it as a skin fold or external artefact. Uncritical acceptance propagates automation bias, causing unnecessary escalation and avoidable patient distress. Professional verification restores clinical proportionality by filtering probabilistic output through the physical conditions of image acquisition [1, 11, 12, 16]. *AI output does not acquire clinical meaning until interpreted in context.*

Accountability Within Regulatory High-Risk Contexts

The EU AI Act categorises many healthcare AI systems as high-risk, imposing requirements for transparency, documentation, traceability, and post-market surveillance [9, 17]. While manufacturers bear primary design responsibility, clinical deployment remains governed by professional duty of care. AI does not assume moral agency: where an algorithm contributes to a suboptimal outcome, accountability remains human — distributed across the radiographer, the clinical team, the institution, and the procurement pathway [9, 16].

AI as Amplifier of Epistemic Authority

The preceding sections identify algorithmic vulnerabilities that epistemic authority must counteract. An equally important dimension concerns the ways in which AI can actively augment the radiographer's supervisory capacity when engaged critically. These framings are not contradictory: epistemic authority is exercised precisely through the decision to accept, modify, or override algorithmic output — and AI, when appropriately governed, extends both the reach and the precision of that decision.

Automated protocol compliance monitoring enables continuous surveillance of acquisition parameters across all examinations, generating standardisation evidence that would be impractical to produce through manual audit alone. In this configuration, the AI system amplifies the radiographer's radiomic stewardship function rather than displacing it. Similarly, AI-assisted anomaly detection in high-throughput workflows expands the proportion of clinically relevant findings that receive professional attention — not because the algorithm resolves the clinical question, but because it extends the percep-

RESEARCH ARTICLE
JAHC | JOURNAL OF ADVANCED HEALTH CARE
OFFICIAL JOURNAL OF TSRM PSTRP

tual scope within which the radiographer operates. Explainability outputs — attention heatmaps, SHAP value distributions, uncertainty confidence scores — when critically interrogated, provide radiographers with structured evidence for professional decision-making unavailable in pre-AI workflows [31, 32]. The radiographer who identifies an implausible attention region, contextualises it against acquisition conditions, and escalates or dismisses accordingly is exercising clinical supervision with greater granularity than was previously achievable. Professional responsibility in this context is not diminished but differentiated: it requires a more sophisticated knowledge base, a more active governance orientation, and a more deliberate engagement with the evidence underlying each AI-mediated output.

The answer to the question *can AI genuinely improve the professional accountability of radiographers?* is therefore affirmative — conditional on the cultivation of the epistemic authority that makes critical engagement possible. AI does not replace professional judgment; under governed conditions, it expands the domain over which that judgment can meaningfully operate.

Radiomics and the Epistemic Stewardship of Data

Radiomics aims to correlate imaging phenotypes with treatment response, prognosis, or biological behaviour through high-dimensional quantitative feature extraction [1, 3, 4]. In practice, radiomic validity is critically dependent on technical reproducibility and acquisition consistency. Radiomic features are acutely sensitive to protocol variation, reconstruction algorithm selection, noise texture, voxel dimensions, and segmentation methodology [1, 3, 20]. A transition from filtered back-projection to iterative reconstruction, for example, alters the noise power spectrum and modifies grey-level texture distributions — changing feature values even when the underlying pathology is biologically unchanged. Where acquisition pipelines lack harmonisation, AI models risk encoding institutional technical signatures rather than biologically meaningful signal, producing associations that appear statistically robust but are epistemically unreliable [1, 20].

Feature Stability and Acquisition Standardisation

Feature stability, evaluated through metrics such as the Intraclass Correlation Coefficient (ICC), is a prerequisite for reliable downstream AI modelling; computational sophistication cannot compensate for epistemically unstable inputs [1, 20]. Published evidence quantifies the magnitude of this vulnerability. Studies examining the impact of reconstruction algorithm transitions on radiomic feature stability have demonstrated that changes from filtered back-projection to iterative reconstruction reduce ICC below the accepted reliability threshold of 0.75 for more than 40% of first-order and texture features [33,34]. Variation in voxel dimensions — for example, a change from 1.25 mm to 2.50 mm slice thickness — can reduce ICC below 0.60 for shape features in lung nodule radiomic pipelines [35]. The commonly applied ICC threshold structure provides a practical reference framework for acquisition standardisation audits: ICC < 0.50 indicates poor reliability, 0.50–0.75 moderate, 0.75–0.90 good, and > 0.90 excellent [36]. These figures illustrate that the gap between what radiomic models appear to measure and what they actually encode is not a hypothetical concern but a quantifiable, protocol-dependent risk whose magnitude is directly addressable through the standardisation practices described in this section.Technical precision in radiographic acquisition is therefore not merely an image-quality obligation — it is a form of epistemic stewardship. Standardisation of kVp, mAs, reconstruction kernel, and slice thickness constitutes simultaneously a technical protocol requirement and a scientific validity condition, positioning radiographers as active gatekeepers of the data provenance upon which AI-derived predictions depend [1, 20, 27].

Performance Metrics and Systemic Vulnerabilities

Numerical performance metrics — AUC-ROC, sensitivity, specificity, Dice coefficient, SSIM — characterise model behaviour under controlled conditions but do not guarantee contextual reliability, population generalisability, or safe workflow integration [1, 3, 12–15, 21]. Reported metrics may be inflated by dataset selection effects, class imbalance, overfitting, or data leakage. Subgroup bias is particularly concerning: models trained on non-representative datasets may systematically underperform for under-served populations, embedding diagnostic inequity within routine clinical infrastructure [14]. For radiographers engaged in governance, performance must be treated as a monitored, dynamic attribute rather than a fixed certification [11, 19, 21, 25].

Automation Bias and Professional Deskilling

Automation bias describes the cognitive tendency to over-rely on algorithmic output, manifesting in radiographic workflows as uncritical acceptance of AI-generated positioning scores, dose recommendations, or quality flags without independent verification [12, 23, 24]. A subtler risk is professional deskilling: as automation absorbs repetitive perceptual tasks, the experiential opportunities through which expert clinical judgement is consolidated become attenuated [3, 5]. This risk is most pronounced among trainees developing practice in AI-saturated environments, acquiring supervisory dispositions before foundational technical competence is established — a developmental inversion critically exposed when tools fail or cases diverge from normative pat-

JAHC
OFFICIAL JOURNAL OF TSRM PSTRP
JOURNAL OF ADVANCED HEALTH CARE
RESEARCH ARTICLE

terns [3, 5, 7]. Mitigation requires curriculum design that preserves unmediated clinical exposure and develops supervisory competence alongside, not in place of, core technical expertise.

Algorithmic Drift and Post-Market Surveillance

Algorithmic drift refers to progressive performance degradation from changes in hardware, software, imaging protocols, or patient demographics [28–30]. Drift is insidious because it is typically gradual, creating latent vulnerabilities that compound silently within routine workflows. Structured post-market surveillance — periodic audits, edge-case review, and escalation mechanisms for recurrent discrepancies — should be regarded as an integral safety function, not an administrative supplement [17, 19, 21, 25, 30]. Radiographers occupy a privileged position for early drift detection given their proximity to AI outputs at the point of acquisition and quality control.

Black-Box Opacity and Explainable AI

Opacity remains a significant safety barrier in clinical AI adoption [1, 16, 21, 31, 32]. Post-hoc explainability methods — SHAP values, LIME, attention heatmaps — do not eliminate opacity but provide a practical supervisory layer, revealing which image regions or input variables drove model output and enabling identification of implausible or artefact-driven algorithmic attention [31, 32]. Explainable AI should be conceptualised as a governance tool: supportive evidence informing human verification, not a substitute for it [21, 22, 31].

Summary for Clinical Practice

The following tables synthesise the principal practice obligations and risk mitigation strategies arising from AI integration in radiographic workflows.

Table 1. *Summary for Clinical Practice in AI-Integrated Radiography*

Domain	Key Professional Obligation
Positioning & quality assurance	Override algorithmic quality flags when repetition is clinically disproportionate (paediatric, frail, trauma patients)
Triage & escalation	Verify AI-flagged findings against acquisition context before escalating; identify out-of-distribution artefacts
Radiomic acquisition	Standardise kVp, mAs, reconstruction kernel, voxel size across protocols; treat consistency as scientific validity condition
Automation bias	Maintain independent clinical verification; do not accept AI outputs without contextual interrogation
Explainable AI	Use attention heatmaps and SHAP outputs as supervisory evidence, not as diagnostic conclusions
Post-market surveillance	Report recurrent AI discrepancies; participate in audit and edge-case review
Governance participation	Engage in procurement decisions, validation processes, and multidisciplinary AI committees

Table 2. *Risk Management and Professional Mitigation Strategies in AI-Integrated Radiography*

Risk	Mechanism	Mitigation Strategy
Automation bias	Uncritical acceptance of AI-generated scores, flags, or recommendations	Structured independent verification at each decision point; Critical AI Literacy training
Professional deskilling	Attenuation of perceptual experience as automation absorbs repetitive tasks	Curriculum design preserving unmediated clinical exposure; intentional non-AI-assisted practice
Algorithmic drift	Gradual performance degradation following hardware, software, protocol, or demographic changes	Periodic audits; longitudinal performance monitoring; escalation pathways for recurrent discrepancies
Dataset bias & inequity	Underperformance on underrepresented populations due to non-representative training data	Subgroup performance assessment at procurement; post-deployment surveillance by patient demographic
Radiomic instability	Acquisition variability producing epistemically unstable feature distributions	Protocol harmonisation; ICC-based feature stability assessment before clinical deployment
Black-box opacity	Inability to interrogate algorithmic reasoning pathway	Require explainability outputs (SHAP, LIME, attention maps) as governance condition at procurement
Regulatory non-compliance	Passive institutional adoption without active professional operationalisation of EU AI Act obligations	Radiographer participation in transparency documentation, traceability, and post-market monitoring

RESEARCH ARTICLE
JAHC JOURNAL OF ADVANCED HEALTH CARE
OFFICIAL JOURNAL OF TSRM PSTRP

JAHC | JOURNAL OF ADVANCED HEALTH CARE
OFFICIAL JOURNAL OF TSRM PSTRP
RESEARCH ARTICLE

Governance, Audit, and Critical AI Literacy

AI governance in radiography is a clinical necessity, not a regulatory formality. The EU AI Act and Medical Device Regulation formally recognise healthcare AI as high-risk, requiring transparency, risk management, human oversight, and post-market monitoring [9, 17, 18, 21, 25]. Regulatory compliance, however, does not automatically translate into safe local practice — institutional governance requires active operationalisation. Effective structures must connect procurement, validation, workflow integration, training, and surveillance, engaging radiographers, radiologists, medical physicists, engineers, IT specialists, and clinical managers [19, 21–25]. AI failures characteristically emerge at technical–clinical–organisational interfaces; governance should be designed around those interfaces, not within isolated professional silos.

Critical AI Literacy as Professional Competency

Critical AI Literacy extends beyond operational familiarity. It constitutes a professional capability encompassing: interrogation of data representativeness; identification of model boundary conditions; recognition of automation bias mechanisms; interpretation of uncertainty outputs; understanding of explainability limitations; and meaningful participation in escalation, audit, and monitoring [11, 16, 21, 22, 25, 27]. It is a prerequisite for responsible use — not a supplementary skill — warranting integration into pre-registration curricula, continuing professional development, and institutional onboarding. Radiographers are evolving from system users into system stewards, with the authority to challenge algorithmic outputs, participate in governance decisions, and advocate for the standardisation and surveillance conditions that make AI clinically safe [11, 20, 22, 27].

Practical Indicators of Effective Human Oversight

The assertion that radiographers must maintain supervisory control over AI outputs requires operationalisation: governance structures must define how effective oversight is measured, not merely that it is expected. Several practical audit indicators can be applied within departmental quality management frameworks:

Override rate (OR): the proportion of AI-generated recommendations — positioning scores, quality flags, triage decisions — that are actively modified or rejected by the operator. A declining OR over time may signal accumulating automation bias rather than genuine AI performance improvement [12, 23].

Correction frequency: the average number of AI output corrections recorded per imaging session, expressed per modality and per operator. Systematic inter-operator variation identifies targeted training needs; temporal variation may indicate algorithmic drift [28–30].

Concordance-with-outcome rate: the proportion of accepted AI recommendations subsequently validated as clinically appropriate in downstream reporting or audit, distinguishing justified acceptance from uncritical deference.

Escalation accuracy: for AI-assisted triage systems, the proportion of AI-flagged cases confirmed as requiring escalation following human verification. Both under-escalation (missed critical findings accepted without correction) and over-escalation (non-significant AI flags propagated without interrogation) constitute auditable failure modes.

These indicators can be integrated into departmental quality dashboards and reviewed on a defined cycle — monthly for high-volume workflows, quarterly for specialised applications. The practical question how often does the operator correct the machine? is not merely operational: it is an institutional governance signal whose longitudinal trend carries direct patient safety implications [21, 25, 30].

Table 3. *Practical Audit Indicators for Human Oversight of AI in Radiographic Workflows*

Audit Indicator	Definition	Safety Signal
Override Rate (OR)	% of AI recommendations modified or rejected	Declining OR ⟶ automation bias risk
Correction Frequency	Mean AI corrections per session (per operator)	Rising trend ⟶ training need; falling ⟶ drift risk
Concordance-with-Outcome Rate	% accepted AI outputs validated downstream	Low rate ⟶ uncritical acceptance pattern
Escalation Accuracy	% AI-flagged cases confirmed on human review	Deviation ⟶ triage reliability failure

Educational and Workforce Implications

The competency demands described throughout this review require direct curricular response. Radiography programmes must move beyond operational AI training to equip practitioners to interrogate, audit, and govern AI systems [11, 16, 21, 25]. Core components should address: AI validation science and performance metric interpretation; bias and representativeness assessment; explainability principles; regulatory literacy; and radiomic standardisation [11, 22, 25, 27]. Continuing professional development should support cognitive bias mitigation, audit methodology, and longitudinal monitoring skills. Hybrid clinical-informatics roles are likely to expand, but curriculum design must equally protect against deskilling by preserving intentional unmediated clinical exposure — supervisory competence must be built upon, not substituted for, foundational technical expertise [3, 5, 7]. A practical implementation model requires structuring AI exposure as a sequential progression rather than simultaneous immersion. In pre-registration curricula, unmediated clinical acquisition — without AI positioning support, automated quality flags, or protocol suggestion engines — should be established as the foundational learning environment, with AI-assisted sessions introduced only after core technical competence has been independently assessed [3, 5, 7]. This developmental sequencing prevents the inversion in which trainees acquire supervisory dispositions before possessing the technical schema against which AI outputs can be meaningfully evaluated. Concrete curriculum components should include: (i) structured error-identification exercises using real or simulated cases in which AI has generated plausible but incorrect outputs, requiring students to locate the failure mechanism; (ii) supervised audit workshops in which trainees apply override-rate tracking and correction-frequency logging to archival datasets; and (iii) protocol reproducibility practicums in which students quantify the feature-level impact of deliberate acquisition parameter variation, directly reinforcing the radiomic stewardship concept at the point of skill formation. Continuing professional development programmes should provide certified pathways in AI audit methodology and clinical governance participation, with competency frameworks aligned to EU AI Act operational requirements [9, 17, 25].

Reframing Professional Identity

AI does not erode radiographic professional identity — it redefines and intensifies it [1, 5, 6, 11, 27]. Professional competence now extends beyond technical execution and image quality optimisation to encompass algorithmic supervision, epistemic stewardship, governance participation, and patient safety advocacy. Radiographers increasingly function as mediators between computational inference and patient-centred care, accountable for clinical quality in environments that are partially but never wholly automatable [1, 11, 16, 20, 21, 27]. This evolution requires institutional recognition, structured professional development, and assertive profession-level participation in procurement decisions, audit processes, multidisciplinary AI committees, and manufacturer discussions on explainability requirements [19, 21–24]. The radiographer who supervises AI is not a diminished professional — they are a more epistemically sophisticated one, accountable for a wider and more consequential range of clinical and governance responsibilities.

Limitations

This review is limited by its narrative design, which does not provide quantitative synthesis and may be subject to selection bias. The evidence base is evolving rapidly; some recommendations will require revision. Empirical studies directly linking AI implementation to radiographer-specific outcomes remain limited, and deskilling trajectories require longitudinal investigation. This work should be interpreted as a practice-oriented synthesis to support professional reflection, curriculum development, and governance planning, rather than a definitive evidence appraisal.

Conclusion

Artificial intelligence transforms radiography through redistribution of responsibility, not displacement of the professional. As procedural tasks become automated, professional judgement becomes simultaneously more visible, more consequential, and less amenable to delegation [1, 11, 16, 21, 27]. AI should be understood as a precision amplifier embedded within complex human–algorithmic healthcare systems. Its clinical value is determined not by computational performance alone, but by the governance structures, standardisation practices, and professional competencies that surround and constrain it. Automation bias, algorithmic drift, radiomic instability, and black-box opacity are not theoretical constructs — they are active, compounding vulnerabilities in real-world radiographic workflows that demand structured, sustained mitigation [12–14, 19–21, 25, 28–32].

The radiographer's role is evolving from technical executor to supervisory clinician and data-quality steward — a transition that represents a genuine expansion of professional authority, not a diminution of it. Sustainable AI integration requires Critical AI Literacy, rigorous technical standardisation, human oversight, and governance architectures capable of detecting and managing systemic risk across the full deployment lifecycle [11, 16, 20, 21, 25].

'The future radiographer is not replaced by AI. The future radiographer governs it.'

RESEARCH ARTICLE
JAHC | JOURNAL OF ADVANCED HEALTH CARE
OFFICIAL JOURNAL OF TSRM PSTRP

JAHC
OFFICIAL JOURNAL OF TSRM PSTRP
JOURNAL OF ADVANCED HEALTH CARE
RESEARCH ARTICLE

References

1. *Malamateniou, C., Hardy, M., Knapp, K. M., & Ramlaul, A. (Eds.). (2026). Artificial Intelligence for Radiographers: Basic Principles, Clinical Applications and Implementation Considerations. Springer. https://doi.org/10.1007/978-3-032-05080-9*
2. *LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436–444.*
3. *Avakian, A., & Bhide, G. B. (2026). Artificial intelligence in radiology: A narrative review of current methods, clinical impact, and future directions. BMC Artificial Intelligence. https://doi.org/10.1186/s44398-025-00020-7*
4. *Alhashim, M., et al. (2025). Artificial intelligence-empowered radiology: Current status and critical review. Diagnostics, 15(3), 282. https://doi.org/10.3390/diagnostics15030282*
5. *Stogiannos, N., Walsh, G., Ohene-Botwe, B., et al. (2025). R-AI-diographers: A European survey on perceived impact of AI on professional identity, careers, and radiographers' roles. Insights into Imaging, 16(1), 43. https://doi.org/10.1186/s13244-025-01918-6*
6. *Walsh, G., Stogiannos, N., Ohene-Botwe, B., et al. (2025). R-AI-diographers: Investigating the perceived impact of AI on radiographers' careers, roles, and professional identity in the UK. Frontiers in Digital Health, 7, 1603511. https://doi.org/10.3389/fdgth.2025.1603511*
7. *Challen, R., et al. (2021). Professionals' responses to the introduction of AI innovations in radiology: A qualitative study. BMC Health Services Research, 21, 541. https://doi.org/10.1186/s12913-021-06861-y*
8. *Aldhafeeri, F. M. (2024). Navigating the ethical landscape of AI in radiography: A cross-sectional study. BMC Medical Ethics, 25, 52. https://doi.org/10.1186/s12910-024-01052-w*
9. *Arvai, N., Katonai, G., & Mesko, B. (2025). Health care professionals' concerns about medical AI: Scoping review. Journal of Medical Internet Research, 27, e66986. https://doi.org/10.2196/66986*
10. *Malamateniou, C., et al. (2021). Artificial intelligence in radiography: Where are we now and what does the future hold? Radiography, 27(Suppl 1), S58–S62.*
11. *Walsh, G., Stogiannos, N., et al. (2023). Responsible AI practice and AI education are central to AI implementation: A rapid review. BJR Open, 5(1), 20230033.*
12. *Tejani, A. S., et al. (2024). Understanding and mitigating bias in imaging AI. RadioGraphics, 44(5), e230067. https://doi.org/10.1148/rg.230067*
13. *Kocak, B., Ponsiglione, A., Stanzione, A., et al. (2025). Bias in artificial intelligence for medical imaging. Diagnostic and Interventional Radiology, 31(2), 75–88. https://doi.org/10.4274/dir.2024.242854*
14. *Seyyed-Kalantari, L., Zhang, H., McDermott, M. B. A., et al. (2021). Underdiagnosis bias of AI algorithms applied to chest radiographs. Nature Medicine, 27(12), 2176–2182.*
15. *Salameh, J. P., et al. (2025). Artificial intelligence for diagnostics in radiology: A rapid systematic scoping review. EClinicalMedicine, 103102. https://doi.org/10.1016/j.eclinm.2025.103102*
16. *World Health Organization. (2021). Ethics and governance of artificial intelligence for health: WHO guidance. WHO.*
17. *European Parliament. (2024). Regulation laying down harmonised rules on artificial intelligence (AI Act). European Union.*
18. *European Society of Radiology. (2022). Current practical experience with AI in clinical radiology: A survey. Insights into Imaging, 13(1), 107. https://doi.org/10.1186/s13244-022-01247-y*
19. *Sujan, M., Smith-Frazer, C., Malamateniou, C., et al. (2023). Validation framework for AI in healthcare: Overview of BS30440. BMJ Health & Care Informatics, 30(1), e100749.*
20. *Ranschaert, E., Rezazade Mehrizi, M. H., Grootjans, W., & Cook, T. S. (2024). AI implementation in radiology: Challenges and opportunities. Springer.*
21. *Brady, A. P., Allen, B., Chong, J., et al. (2024). Developing, purchasing, implementing and monitoring AI tools in radiology. Journal of Medical Imaging and Radiation Oncology, 68(1), 7–26. https://doi.org/10.1111/1754-9485.13612*
22. *Stogiannos, N., Malik, R., Kumar, A., et al. (2023). Black box no more: AI governance frameworks for medical imaging in the UK. British Journal of Radiology, 96(1152), 20221157.*
23. *Stogiannos, N., Litosseliti, L., O'Regan, T., et al. (2024). AI governance and adoption in medical imaging: A cross-sectional survey. International Journal of Medical Informatics, 186, 105423.*
24. *Stogiannos, N., Gillan, C., Precht, H., et al. (2024). A multidisciplinary approach for AI implementation in medical imaging. Journal of Medical Imaging and Radiation Sciences, 55(4), 101717.*
25. *Kotter, E., D'Antonoli, T. A., Cuocolo, R., et al. (2025). ESR recommendations for implementation of the European AI Act. Insights into Imaging, 16(1), 33.*
26. *Hardy, M., & Harvey, H. (2020). Artificial intelligence in diagnostic imaging: Impact on the radiography profession. British Journal of Radiology, 93(1108), 20190840.*
27. *Malamateniou, C., O'Regan, T., McFadden, S. L., & Jackson, M. (2024). AI in radiography practice, research and education: A contemporary review. Radiography, 30(Suppl 2), 56–59.*

28. Vela, D., Sharp, A., Zhang, R., et al. (2022). *Temporal quality degradation in AI models. Scientific Reports*, 12, 11654.
29. de Vries, C., Colosimo, S., Staff, R., et al. (2023). *Impact of different mammography systems on AI performance in breast cancer screening. Radiology: Artificial Intelligence*, 5(3), e220146. https://doi.org/10.1148/ryai.220146
30. Feng, J., Phillips, R. V., Malenica, I., et al. (2022). *Clinical AI quality improvement: Towards continual monitoring and updating of AI algorithms. npj Digital Medicine*, 5, 66.
31. Amann, J., Blasimme, A., Vayena, E., et al. (2020). *Explainability for AI in healthcare: A multidisciplinary perspective. BMC Medical Informatics and Decision Making*, 20, 310.
32. Champendal, M., Muller, H., Prior, J. O., & dos Reis, C. S. (2023). *Interpretability and explainability of AI in medical imaging: A scoping review. European Journal of Radiology*, 169, 111159.
33. Shafiq-ul-Hassan, M., Zhang, G. G., Latifi, K., et al. (2017). *Intrinsic dependencies of CT radiomic features on voxel size and number of gray levels. Medical Physics*, 44(3), 1050–1062. https://doi.org/10.1002/mp.12123
34. Mackin, D., Fave, X., Zhang, L., et al. (2015). *Measuring computed tomography scanner variability of radiomics features. Investigative Radiology*, 50(11), 757–765. https://doi.org/10.1097/RLI.0000000000000180
35. van Timmeren, J. E., Leijenaar, R. T. H., van Elmpt, W., et al. (2016). *Test-retest data for radiomics feature stability analysis: Generalizable or study-specific? Tomography*, 2(4), 361–365. https://doi.org/10.18383/j.tom.2016.00208
36. Koo, T. K., & Mae, M. Y. (2016). *A guideline of selecting and reporting intraclass correlation coefficients for reliability research. Journal of Chiropractic Medicine*, 15(2), 155–163. https://doi.org/10.1016/j.jcm.2016.02.012

Institutional Review Board Statement: Not applicable. This is a narrative review of published literature; no new human or animal data were col lected.

Data Availability Statement: Not applicable. All data discussed in this review are available in the cited peer-reviewedpublica-tions.

Acknowledgments: The author thanks the reviewers for their constructive feedback, which substantially improved the metho dological transparency and clinical utility of this manuscript.

Conflicts of Interest: The author declares no conflict of interest.

RESEARCH ARTICLE