



The medical physicist's role in the artificial intelligence era: an updated EFOMP curricular and professional framework[☆]

George C. Kagadis^{a,*}, Oliver Diaz^{b,c}, Richard Meades^d, Ruben Pauwels^e,
Niall Colgan^f, Federica Zanca^g, Eleftherios Tzanis^h, Dimitrios Visvikisⁱ

^a 3DMI Research Group, Department of Medical Physics, School of Medicine, University of Patras, Rion, Greece

^b Departament de Matemàtiques i Informàtica, Universitat de Barcelona, Barcelona, Spain

^c Computer Vision Center, Universitat Autònoma de Barcelona, Bellaterra, Spain

^d Royal Free London NHS Foundation Trust, Royal Free Hospital, Nuclear Medicine Department, Pond Street, London NW3 2QG, United Kingdom

^e Department of Dentistry and Oral Health, Aarhus University, Aarhus, Denmark

^f Department of Physics, School of Natural Sciences, College of Science and Engineering, University of Galway H91 TK33 Galway, Ireland

^g European Innovation Council and SMEs Executive Agency (EISMEA) Bd. Simon Bolivar 34, the North Light Building (SB34), 3rd floor, 1000 Brussels, Belgium

^h Artificial Intelligence and Translational Imaging (ATI) Lab, Department of Radiology, School of Medicine, University of Crete, Heraklion, Crete, Greece

ⁱ LaTIM, UMR 1101, Inserm, University of Brest, Brest, France

ARTICLE INFO

Keywords:

Artificial intelligence

Education

Medical physics expert

ABSTRACT

Purpose: Five years after the European Federation of Organizations for Medical Physics (EFOMP) published its first curricular framework for artificial intelligence (AI) in medical physics, the updated core curricula for radiology, radiotherapy and nuclear medicine have incorporated AI as a core competency rather than an addendum, a change consistent with the framework's principal recommendation. This paper examines the resulting educational landscape and defines the next stage of AI competency for the profession.

Materials and Methods: Focus Group 3 of the EFOMP Artificial Intelligence Special Interest Group used a structured, iterative process of expert deliberation combining comparative curriculum analysis, competency mapping, regulatory analysis and iterative expert deliberation within the Working Group. Four primary source documents were analyzed in depth: the 2021 EFOMP white paper and the three updated EFOMP subspecialty core curricula for radiology, radiotherapy and nuclear medicine. The AI-related content of the three curricula was systematically compared against a predefined analytical framework covering scope, depth of treatment, European Credit Transfer and Accumulation System (ECTS) allocation, terminology, regulatory references, and treatment of recent AI developments. Competency gaps beyond the entry-level curricula were identified and articulated in the knowledge–skills–competence format used in EFOMP curricula. The resulting framework was informed by three large-scale surveys of medical physicists conducted in 2020, 2024 and 2025, and underwent iterative review within the Working Group before consultation within the EFOMP community.

Results: The three updated curricula converge on a shared methodological foundation but diverge in ECTS allocation, coverage of recent AI architectures, and treatment of the EU AI Act, and six competency gaps were identified as common to all three subspecialties. A cross-subspecialty advanced competency framework is proposed, organized into four domains: validation and quality assurance, AI lifecycle management, generative AI and foundation models, and governance. A professional lifecycle model is introduced describing how these competencies should be acquired during specialist training, maintained through accredited continuing professional development, and formally recognized throughout a medical physicist's career. The medical physicist's responsibilities under the EU AI Act are defined, and the proposed framework is translated into practical recommendations for EFOMP, its National Member Organizations, training centers, and individual practitioners.

Conclusions: The profession's challenge has evolved from demonstrating the importance of AI to ensuring its safe, trustworthy, effective and regulatory-compliant integration into clinical practice, thereby strengthening patient safety and quality of care.

[☆] EFOMP AI Special Interest Group, Focus Group 3.

* Corresponding author.

E-mail addresses: kagadis@upatras.gr, gkagad@gmail.com (G.C. Kagadis).

1. Introduction and context

Five years have passed since EFOMP published the first curricular and professional framework for artificial intelligence (AI) in medical physics [1]. That document, prepared by the EFOMP Working Group on AI in Medical Physics, proposed a two-level Basic and Advanced curriculum to be read as an addendum to the then-existing EFOMP subspecialty core curricula for medical physicists in radiology (2011), radiotherapy (2011) and nuclear medicine (2013). It was, by the explicit acknowledgement of its authors, an interim measure: a response to a clear need at a moment when the subspecialty core curricula did not yet reflect the arrival of AI in clinical practice.

The landscape has shifted on three fronts since 2021.

First, the **AI technology** itself has changed substantially. The 2021 curriculum was framed around convolutional neural networks, classical machine learning models, and the integration of commercial AI software-as-medical-device into clinical workflows. Since then, generative AI (including large language models), foundation models and agentic AI have moved from research curiosity toward clinically relevant applications across medical imaging and healthcare, fundamentally changing the AI landscape in medicine [2,3].

Second, the **regulatory landscape** has been transformed. The EU AI Act entered into force in 2024 and classifies most clinical AI systems as high-risk. Following the political agreement on the “Digital Omnibus on AI” (7 May 2026), the rules for high-risk AI systems embedded in regulated products have an extended transition period until 2 August 2028, with standalone high-risk systems applying from 2 December 2027. Its requirements for risk management, data governance, transparency, human oversight, accuracy and post-market monitoring intersect substantially with the Medical Device Regulation (MDR) and with tasks that the 2013/59/Euratom Basic Safety Standards Directive and EFOMP Policy Statement 16 already place within the remit of the Medical Physics Expert (MPE). Throughout this paper, “medical physicist” and “MPE” are used interchangeably; the latter denotes the specific regulatory function defined under 2013/59/Euratom and EFOMP Policy Statement 16, following EFOMP convention. The MPE is now positioned, by the convergence of these instruments, as a documented contributor to the safe and effective deployment of clinical AI, a role that the 2021 paper anticipated but could not yet describe in regulatory terms.

Third, and most consequentially for this document, the **EFOMP curricular landscape** has been updated. New subspecialty core curricula have been published for medical physics experts in radiology, radiotherapy and nuclear medicine, each of which now contains a dedicated section on data processing, statistics, modeling and AI. These sections — different in depth, scope, terminology and ECTS allocation, but unified in their recognition that AI is now a core competency rather than an addendum — have absorbed much of what the 2021 Basic level was designed to provide. The premise on which the original white paper rested is therefore no longer the operative gap.

These developments fundamentally change the educational challenge. AI is no longer an emerging technology that can be appended to an otherwise stable curriculum. It is now embedded across all medical physics subspecialties, evolves faster than conventional curriculum revision cycles, and operates within a rapidly developing regulatory environment. The question is therefore no longer whether AI should be taught, but how AI competency should be structured, maintained and recognized throughout professional practice.

The framework presented here is additionally informed by ongoing work within the EFOMP AI Special Interest Group (AI SIG), including a European survey of medical physicists' priorities for AI competencies and clinical implementation.

This paper proposes a new framework for AI competency in medical physics that reflects the technological, educational and regulatory developments over the past five years. Rather than revisiting the foundational AI concepts that are now embedded within the EFOMP

subspecialty curricula, it focuses on the competencies required to safely implement, evaluate, govern and continuously adapt AI in clinical practice. The framework extends beyond specialist training to encompass advanced and expert competencies, lifelong professional development, and the evolving responsibilities of the Medical Physics Expert within the European regulatory landscape.

In doing so, this white paper updates and supersedes Zanca et al. [1]. It builds directly on that document and treats it as the foundation on which a more mature framework can now stand. The 2021 paper made the case that medical physicists *must engage with AI*; that case is now settled. The task ahead is to define *how* that engagement is structured, sustained and recognized across European medical physics practice.

2. Methodology

This white paper was developed by the EFOMP AI Special Interest Group, Focus Group 3 using a structured consensus methodology combining comparative curriculum analysis, competency mapping, regulatory analysis and expert consensus. The objective was to update the 2021 EFOMP AI competency framework to reflect technological, educational and regulatory developments since its publication.

The work proceeded in four steps.

Literature review. The focus group identified four primary source documents to be analyzed in depth: the 2021 white paper [1], and the three updated EFOMP subspecialty core curricula for medical physics experts in radiology [4], radiotherapy [5], and nuclear medicine [6], each of which now contains a dedicated section on data processing, statistics, modeling and AI. Secondary source documents included EFOMP Policy Statement 16 on the role and competences of medical physicists and medical physics experts under 2013/59/Euratom, the European Guidelines on Medical Physics Expert (Radiation Protection 174), recent EFOMP focus group outputs and position statements, and relevant statements from sister organizations (notably ESR, EANM and ESTRO) where these inform the regulatory or clinical context.

Comparative curriculum mapping. The AI-related content of the three subspecialty curricula was systematically compared using a predefined analytical framework covering scope of AI topics, depth of treatment, ECTS allocation, terminology, regulatory references, and explicit treatment of recent developments (generative AI, foundation models, the EU AI Act). This mapping, presented in Section 3 and the appendices, made visible the convergences, divergences and gaps that any harmonized framework must address. The mapping was carried out by the focus group's lead reviewers against the predefined framework and cross-checked collectively by the full group at subsequent working meetings; no independent-coder agreement statistic was calculated. It is accordingly best characterized as a structured expert narrative synthesis of the published curricula rather than a systematic content analysis with a formally validated coding protocol, and the framework is presented on that basis throughout this paper.

Competency Gap analysis. The focus group identified competency areas that fall outside or beyond the entry-level subspecialty curricula — notably advanced and expert-level competencies in AI lifecycle management, generative AI, foundation models, and AI governance under the EU AI Act. These were articulated in the knowledge–skills–competence (KSC) format used in EFOMP curricula, ensuring methodological continuity with the existing curricular framework and enabling direct integration into future curriculum revisions.

Consensus development. The resulting framework underwent iterative review within the Working Group and was subsequently circulated for consultation within the EFOMP community before finalization. This process did not follow a named, formal consensus methodology (e.g., a Delphi process, nominal group technique, or the RAND/UCLA appropriateness method) with defined rounds, structured anonymity or a quantified agreement threshold. “Consensus” in this paper therefore denotes agreement reached through the Working Group's iterative deliberation and EFOMP community consultation, and should

be read as such rather than as a claim to formally quantified professional consensus.

Accordingly, the proposed framework should be regarded as the focus group's expert judgment, reached through iterative deliberation rather than a formally quantified consensus method, informed by comparative curriculum analysis and the current European regulatory landscape, rather than as a systematic review of the literature. Because AI technologies and their governance continue to evolve rapidly, the framework incorporates a lifelong professional development perspective to support future adaptation beyond periodic curriculum revision.

Sections 3 and 4 present the results of this analysis: Section 3 reports the comparative curriculum mapping — convergences, divergences and gaps — while Section 4 sets out the advanced and expert competency framework built on that analysis, before Sections 5–7 place the framework in a professional-lifecycle, regulatory and implementation context.

3. Harmonization across the three EFOMP subspecialty curricula

The most consequential change since the 2021 white paper is that the EFOMP subspecialty core curricula have been updated, and each now contains a dedicated section on data processing, statistics, modeling and artificial intelligence. The radiology curriculum addresses these topics in Section III.6 with an allocation of 10 European Credit Transfer and Accumulation System (ECTS); the radiotherapy curriculum in Section III.9 with 8 ECTS; and the nuclear medicine curriculum in Section III.12 with 15 ECTS. AI has therefore moved, in EFOMP doctrine, from addendum to core. This represents a significant maturation of the curricular framework and validates the trajectory established by the 2021 paper.

The harmonization question now applies to the coherence between the three subspecialties. A practicing MPE, whether in radiology, radiotherapy or nuclear medicine, should share a common foundation in AI competency, with subspecialty-specific extensions layered on top. The focus group analyzed the three curricula against this benchmark. The findings are presented below in three parts: convergences (where the three curricula already align), divergences (where they differ in ways that warrant attention), and gaps (where all three are silent on competencies the focus group considers essential).

3.1. Convergences

The three curricula align substantially on the core AI competencies expected of an MPE.

All three articulate the MPE's role in identical strategic terms: as a bridge between clinical professionals and manufacturers of AI-based systems, with primary responsibility for ensuring that AI tools are used safely and effectively, in line with current regulations including General Data Protection Regulation (GDPR) and MDR. All three describe the MPE as responsible for local validation and configuration of AI tools, and all three open the door to MPE involvement in clinical innovation, research and the development of new AI applications. This convergence on professional role is itself a significant harmonization outcome and provides a stable foundation for the rest of the framework.

All three curricula share the same core statistical and methodological foundation:

- **Statistical foundations** — probability, statistical inference, and study design
- **Machine learning models** — the same core families across all three curricula
- **Methodological concerns** — bias-variance trade-off, overfitting, cross-validation, regularization
- **Deep learning** — basics of deep learning and convolutional neural networks

- **QA of AI/ML** — listed as an expert competency in all three, signaling shared recognition that quality assurance of AI is now a defining MPE responsibility

The harmonized core that emerges from these convergences should be regarded as the *de facto* common foundation of EFOMP AI training, and any further EFOMP guidance.

3.2. Divergences

Beyond the harmonized core, the three curricula diverge in ways that have practical consequences for trainees, training centers and the recognition of MPE qualifications across European jurisdictions.

ECTS allocation. Each of the three subspecialty curricula totals 240 ECTS; within that total, the 15 ECTS assigned to AI/data content by the nuclear medicine curriculum is nearly double the 8 ECTS allocated by radiotherapy, with radiology at 10 ECTS in between. Some of this difference reflects genuine variation in the role AI plays in each subspecialty — the deep integration of AI into PET reconstruction, attenuation correction and dosimetry in nuclear medicine warrants substantial training time. The published curricula do not document the basis for these allocations, so it is not possible to determine how much of the difference is principled, and the focus group recommends that the next round of curriculum revisions revisit ECTS allocations against a common reference framework rather than independently.

Coverage of recent AI architectures. Only the nuclear medicine curriculum explicitly names transformers, foundation models, generative adversarial networks (GANs) and recurrent neural networks among the AI architectures trainees should be aware of. The radiology and radiotherapy curricula remain at the level of CNNs and classical machine learning models. Given that transformer-based architectures and foundation models are now embedded in commercial AI systems across all three subspecialties — including radiology image reconstruction and reporting, and radiotherapy auto-contouring — this asymmetry should be corrected at the next revision. The focus group recommends that a common minimum coverage of architecture families be specified across all three subspecialty curricula.

Treatment of the EU AI Act and AI ethics. The radiology curriculum is alone among the three in explicitly addressing the EU AI Act, citing the ESR's recommendations for its implementation [7] and the joint European–North American multisociety statement on the ethics of AI in radiology [8]. Radiotherapy and nuclear medicine reference GDPR and MDR but do not cover the AI Act or AI-specific ethics. This is the most consequential divergence the focus group identified, because the AI Act applies across all clinical AI systems, regardless of subspecialty. Section 6 of this paper addresses the AI Act in detail; in curricular terms, the focus group recommends that the AI Act and AI ethics be incorporated as common content across all three subspecialty curricula at next revision.

Subspecialty-specific applications. The three curricula appropriately differ in the AI applications they emphasize. Radiology highlights AI-assisted scan planning, acquisition parameter optimization, image reconstruction and post-processing, segmentation, lesion detection, and dosimetry. Radiotherapy emphasizes pre-treatment QA, conversion of CBCT pixel values into Hounsfield units, pseudo-CT generation from MRI, treatment planning, contouring and treatment outcome prediction. Nuclear medicine focuses on AI in image reconstruction, tumor volume segmentation, pseudo-CT generation for PET/MR attenuation correction, low-dose imaging, predictive modeling and radiomics. This divergence is appropriate and reflects genuine differences in clinical practice. It should be preserved at the subspecialty-specific layer of training while the methodological foundation is harmonized.

3.3. Gaps

Several competency areas are insufficiently addressed, or absent,

across all three subspecialty curricula. The focus group identifies six such gaps that this white paper aims to address through the advanced and expert competency layer (Section 4) and the regulatory framework discussion (Section 6). A candidate topic was counted as a gap, rather than simply an underdeveloped area already present in the curricula, only where it met three criteria agreed by the focus group: the topic was absent, not merely underdeveloped, from all three subspecialty curricula; it was already generating clinical or operational AI applications at the time of writing; and it carried a plausible safety, quality-assurance or regulatory implication for the MPE. The six gaps below are those that met all three criteria in the focus group's review; the specific number reflects this criterion-based screening rather than a target count.

The first gap is **AI lifecycle management** beyond procurement and commissioning. The curricula address the input side of the lifecycle (procurement, validation, integration) but say comparatively little about ongoing performance monitoring, drift detection, model retraining decisions, or decommissioning of AI systems whose performance has degraded. This evolution is also reflected in emerging international guidance advocating continuous post-deployment monitoring of AI systems rather than one-time validation [9].

The second gap is **generative AI and foundation models** in clinical and operational workflows, including large language models (LLMs). None of the three curricula meaningfully address the safety, validation and quality assurance challenges raised by these systems — whether in image synthesis, report generation, clinical decision support, scheduling or patient communication — even though they are entering practice at pace and are already embedded across European healthcare systems. The MPE's role in validating and assuring such systems is essentially absent from the curricula.

The third gap is **agentic AI systems**. Recent advances have moved beyond standalone LLMs toward autonomous or semi-autonomous multi-agent systems that combine foundation models with planning, memory, tool use, API/function calling and multi-step decision-making, capable of automating complex clinical and operational workflows including report generation, protocol optimization, AI model development, validation, deployment, performance monitoring and quality assurance [10–12]. As these systems become more integrated into clinical and research environments, their governance, safety and quality assurance should be recognized as an advanced competency for the Medical Physics Expert; the curricula do not yet address them.

The fourth gap is **operationalizing human oversight under the EU AI Act**. The Act requires meaningful human oversight of high-risk AI systems, but the curricula do not yet articulate what this means in MPE practice — what the MPE documents, how oversight is exercised, or how it is audited.

The fifth gap is **post-market surveillance methodology specific to AI**. The curricula reference post-market surveillance generically (as part of MDR), but AI systems present surveillance challenges — performance drift, distributional shift, silent degradation — that conventional medical device surveillance frameworks do not fully address.

The sixth gap is **AI-related cybersecurity**. The curricula address general medical device cybersecurity but not the AI-specific attack surface. Adversarial inputs, model extraction, training data contamination are increasingly relevant as AI systems become a larger proportion of clinical infrastructure.

These gaps are not failings of the subspecialty curricula, which were necessarily framed around topics with established teaching material. They are, rather, the natural frontier of MPE competency in a rapidly evolving field, and they define the agenda for Sections 4 through 6 of this paper.

3.4. Summary

The three updated EFOMP subspecialty curricula represent a substantial advance over the 2021 baseline: they have established AI as core

competency, converged on a shared methodological foundation, and articulated a coherent professional role for the MPE in AI deployment. The divergences identified above are tractable and should be addressed at the next round of curriculum revisions; the gaps require a complementary framework of advanced, expert and lifelong learning competencies, which the remainder of this white paper sets out. A detailed comparative mapping of the three curricula against the analytical framework is provided in Appendix A.

4. The advanced and expert competency layer

The six gaps identified in Section 3.3 do not divide neatly among the three subspecialties: they are largely subspecialty-agnostic, arising from the pace of AI development rather than from any clinical domain. The focus group's position is that they are best addressed not by further subdividing the subspecialty curricula, but by a common advanced and expert competency layer that sits above the harmonized entry-level foundation identified in Section 3.1. The four domains that structure this layer (Section 4.3) are defined broadly enough to apply across radiology, radiotherapy and nuclear medicine alike; the specific competencies within each domain, however, are defined at the subspecialty level, reflecting the different AI applications, risk profiles and clinical workflows of each.

4.1. Empirical basis: three surveys, 2020–2025

The case for this competency layer rests on a consistent, five-year body of survey evidence rather than on the focus group's judgment alone. Three large-scale surveys, conducted in 2020, 2024 and 2025, together provide a longitudinal picture of AI adoption and preparedness within the medical physics profession [13–15].

The 2020 EFOMP survey established the baseline: 81% of respondents believed AI would improve their daily work, but self-reported AI knowledge averaged only 2.3 out of 5, nearly two-thirds were not involved in any AI-related project, and 88% identified a clear need for dedicated AI training [13]. The 2024 UK IPPEM survey found AI by then embedded in routine practice — 59% reported regular use, particularly in radiotherapy — yet only 10% had received formal AI education and only 15% considered themselves prepared. Those with prior formal education were roughly three times more likely to feel prepared, and insufficient education, expertise and governance were the principal barriers to safe deployment [14]. The 2025 EFOMP AI SIG survey shows continued maturation but persistent gaps: 71% reported using clinical AI tools, yet only 45% reported direct MPE involvement in implementation, 40% reported no previous AI education at all, and only 11% considered their employer to be adequately preparing the workforce. As before, structured AI education was the strongest predictor of preparedness, and a national AI initiative was significantly associated with higher adoption, involvement and preparedness [15]. The apparent gap between the 10% reporting formal AI education in 2024 and the 40% reporting no previous AI education at all in 2025 is likely attributable to the differently worded survey questions (receipt of formal education specifically, versus any prior education at all) rather than to a genuine reversal in the underlying trend.

Read together, these three surveys differ in sponsoring body, geographic scope, sample size and question wording and are not directly comparable; read as a set of related snapshots rather than a single continuous trajectory, they suggest a profession in which technology adoption has consistently outpaced structured education, governance and quality assurance capacity. Together these findings motivate the competency layer set out below: adoption has risen steadily while structured education, governance and quality-assurance capacity have not kept pace.

4.2. Topics of interest for AI training

The 2025 AI SIG survey also ranked interest in specific training topics: (1) validation of AI performance; (2) medical image analysis with AI; (3) the role of the medical physicist in AI; (4) big data analysis; (5) basics of machine and deep learning; (6) regulatory issues in AI; (7) advanced machine and deep learning; and (8) ethical issues in AI [15].

A Jaccard similarity analysis of co-selected topics — reported in [15], a manuscript still in preparation at the time of writing, for which the clustering threshold, algorithm, validation approach, sample size, response rate and the statistical definition used for the country-level comparison below are not yet publicly reported, and this finding should accordingly be treated as provisional pending that publication — identified in three clusters: Cluster 1 (validation, image analysis, MPE role) links applied technical assurance with professional identity; Cluster 2 (regulatory and ethical issues) reflects a distinct governance concern; Cluster 3 (regulatory issues and validation) bridges the two. Only interest in “the role of the medical physicist in AI” varied significantly by country, suggesting most training priorities are shared broadly across Europe [15].

This ranking directly shaped the domain prioritization below: validation (Domain 1) and the MPE’s governance role (Domain 4) sit at the top of both the ranked list and the strongest cluster, ahead of the more specialized content of generative AI (Domain 3) — though all four domains are needed to close the gaps identified in Section 3.3.

4.3. The four competency domains

The advanced layer is organized into four domains, each mapped to one or more of the six gaps and grounded in the survey evidence above. Each domain is defined using the three-level competence framework (expert, collaborative, awareness) already established in the subspecialty curricula, so it integrates into existing training and assessment structures without a parallel framework.

Domain 1: Advanced validation and quality assurance (gap 6; technical foundation for gap 5; the top-ranked training topic in Section 4.2) extends existing QA competencies to the specific challenges of validating AI systems: performance evaluation across representative and edge-case populations, calibration and uncertainty assessment beyond accuracy alone, and the AI-specific attack surface — adversarial inputs, model extraction, training data poisoning — that conventional device security testing does not address.

- **Expert** competence means designing and executing such a validation protocol independently
- **Collaborative** competence means contributing technical input under another expert’s leadership
- **Awareness** competence means understanding why AI validation differs from conventional device QA

Domain 2: AI lifecycle management (gap 1; combined with Domain 1, gap 5) covers what happens after initial validation: ongoing performance monitoring, detection of drift and silent degradation, criteria for revalidation or retraining, and structured decommissioning.

- **Expert** competence means designing and operating such a monitoring program
- **Collaborative** competence means operating one designed by another expert
- **Awareness** competence means recognizing that AI performance is not static and requires distinct, ongoing responsibility

Domain 3: Generative AI, foundation models, and agentic AI (gaps 2 and 3) cover the safety, validation and governance of generative models — image synthesis, report generation, large language models, and agentic AI systems — including their distinct failure modes:

hallucination, bias inherited from foundation-scale pretraining, planning and tool-use errors, unintended autonomous actions, and the difficulty of exhaustive testing given open-ended outputs.

- **Expert** competence means designing validation approaches appropriate to these outputs and advising on safe integration
- **Collaborative** competence means contributing under another expert’s leadership
- **Awareness** competence means understanding that generative systems fail differently from discriminative models and cannot be assured with the same test paradigms

Domain 4: AI governance (gap 4; among the top training priorities in Section 4.2; developed further in Section 6) covers the operational translation of regulatory human-oversight requirements into departmental practice — what oversight looks like for a given system, what audit trail it generates, and how the MPE participates in institutional governance.

- **Expert** competence means translating the EU AI Act’s oversight requirement into a specific, auditable protocol and leading governance structures
- **Collaborative** competence means participating under another expert’s leadership
- **Awareness** competence means understanding the existence and purpose of the Act’s oversight, transparency and post-market monitoring obligations

These four domains are not a second full curriculum to be delivered to every MPE at qualification. They define a competency ceiling the profession as a whole must reach, distributed across a career through the lifecycle model in Section 5. Not every MPE will develop expert-level competency in all four; consistent with EFOMP’s existing model for other cross-cutting competencies, departments and networks are expected to maintain this competency collectively, with individuals developing depth where most relevant to their setting.

Fig. 1 illustrates the gap-to-domain mapping schematically; the full mapping is tabulated in Appendix B, and the survey-priority mapping in Appendix C.

5. AI competencies across the MPE professional lifecycle

A framework of competencies is incomplete without a corresponding framework for how those competencies are acquired, maintained and recognized across the professional lifecycle. This section sets out that lifecycle model in four stages: initial specialist training, transition to independent practice, continuing professional development, and advanced or expert recognition. Fig. 2 summarizes this model and the recurrent nature of the continuing professional development (CPD) stage.

At the initial specialist training stage, AI competency is acquired as part of the subspecialty core curriculum. The harmonized entry-level foundation identified in Section 3.1 — the statistical and methodological basis of machine learning, the principles of clinical evaluation, and the core regulatory context — belongs here, delivered within the structured, assessed training programs that lead to recognition as an MPE. The focus group’s position is that this material is now core rather than elective, and that trainees should not complete specialist training without demonstrated competency in it. This is a change from the 2021 framework, which treated AI as an addendum to be layered onto existing curricula [1].

At the transition to independent practice, the newly qualified MPE consolidates entry-level AI competency in the clinical setting and begins to encounter the advanced competencies of Section 4 in context — participating in the validation of a clinical AI system, contributing to its ongoing monitoring, or supporting departmental governance. The focus




Fig. 1. The advanced and expert competency layer: mapping of the six curricular gaps (Section 3.3) to the four competency domains (Section 4.3). Domains 1 and 3 each absorb two gaps; Domains 2 and 4 each absorb one. Domain 1 also underpins gap 5 (post-market surveillance) together with Domain 2. See Appendix B for the full mapping in tabular form.

AI competencies across the MPE professional lifecycle




Fig. 2. AI competencies across the MPE professional lifecycle: four stages from initial training to expert recognition. Initial training, transition and continuing professional development (CPD) form a sequential path; CPD is recurrent throughout the remainder of the career rather than a single event. Advanced or expert recognition is an optional branch available to MPEs who develop competency beyond the general-practitioner level, not a fifth sequential stage.

group recommends that this transition be structured rather than incidental, with the advanced competency domains introduced through supervised practice and departmental mentorship rather than left to individual initiative.

At the continuing professional development stage, which spans the majority of a career, the emphasis shifts from acquisition to maintenance and updating. The focus group proposes that CPD for AI competency be organized around three principles. First, it must be recurrent: because the field changes continuously, AI CPD cannot be a single accredited course but must be a sustained program, and the focus group recommends that AI-specific CPD carry a minimum recommended allocation within the MPE's overall CPD portfolio. Second, it must be structured around the four advanced competency domains of [Section 4](#), so that CPD activity maps onto a coherent competency framework rather than accumulating as disconnected courses. Third, it must be accredited: CPD activities should be assessed and recognized by EFOMP or its NMOs against the competency framework, so that a departmental or regulatory body can rely on the currency of an MPE's AI competency. The absence of such accreditation is, at present, a significant gap.

This CPD emphasis is directly supported by the survey evidence presented in [Section 4.1](#): across all three surveys spanning 2020 to 2025, structured formal education was consistently the strongest predictor of preparedness, while informal or self-guided learning — the dominant mode reported by respondents — left many only partly prepared [[13–15](#)]. A recurrent, accredited CPD pathway is the profession's structural answer to this five-year pattern.

At the advanced or expert recognition stage, the MPE who has developed competency across the four domains to a level beyond that expected of the general practitioner may seek formal recognition of advanced or expert status in AI. The focus group recommends that EFOMP establish such a recognition mechanism, analogous to existing routes for expert recognition in other domains, so that departments, national authorities and employers can identify the MPEs qualified to take professional responsibility for the most demanding AI work — the validation of high-risk systems, the design of surveillance programs, and the technical leadership of institutional AI governance. [Section 7](#) addresses how this mechanism should be established.

6. The MPE's role under the EU AI Act and evolving regulation

The competency framework set out in the preceding sections is not developed in a regulatory vacuum. Its urgency and much of its specific content derive from a legal instrument that did not exist when the 2021 framework was written: Regulation (EU) 2024/1689, the Artificial Intelligence Act, which entered into force in 2024 and is being phased into application over the following years [[16](#)]. The timeline most relevant to the MPE concerns high-risk AI systems embedded in regulated products — including medical devices and in vitro diagnostic medical devices (IVDs) under the MDR and the In Vitro Diagnostic Regulation (IVDR), the category that covers the great majority of clinical AI. Following the political agreement reached on the “Digital Omnibus on AI” on 7 May 2026, full high-risk obligations for these systems now apply from 2 August 2028, a year later than the Act's original 2 August 2027 date; stand-alone high-risk systems under Annex III apply from 2 December 2027. Together with the Medical Device Regulation (MDR), the GDPR, the Basic Safety Standards Directive (2013/59/Euratom) and EFOMP Policy Statement 16, the AI Act defines a regulatory environment in which the MPE has a distinct and, in several respects, newly formalized role. This section sets out that role.

The AI Act classifies AI systems by risk. Most clinical AI systems — those intended to inform diagnosis, guide treatment or otherwise support clinical decisions — fall within the high-risk category, because they are either themselves medical devices or safety components of medical devices already regulated under the MDR. High-risk classification triggers a substantial set of obligations: risk management across the system lifecycle, data governance and quality requirements for training and

testing data, technical documentation, record-keeping and traceability, transparency toward users, human oversight, and appropriate levels of accuracy, robustness and cybersecurity. Each of these obligations maps onto competencies developed in this paper, and several map directly onto the established remit of the MPE.

The requirement for human oversight is central to the MPE's role. The AI Act requires that high-risk AI systems be designed so that they can be effectively overseen by natural persons, with the aim of preventing or minimizing risks to health and safety. In the clinical setting, giving operational meaning to this requirement — determining what oversight of a specific AI system looks like in practice, what the overseeing clinician must be able to see and do, and how oversight is documented and audited — is a technical task that falls naturally to the MPE. The focus group's position is that the MPE is the professional best placed to translate the Act's abstract oversight requirement into concrete, validated clinical practice, and that this responsibility should be explicitly recognized in national implementations. This is precisely the competency developed in Domain 4 of [Section 4.3](#).

The requirements for accuracy, robustness and post-market monitoring likewise fall within the MPE's competency. The validation and lifecycle-management domains of [Section 4](#) (Domains 1 and 2) are, in substance, the technical means by which a department demonstrates compliance with the Act's requirement that high-risk systems achieve and maintain appropriate performance across their lifecycle. The MPE who validates a system before deployment, monitors it for drift, and determines when it must be revalidated or withdrawn is performing precisely the assurance function the Act requires. This is the strongest single argument for treating the advanced competency layer as core rather than optional: it is the profession's contribution to a legal obligation that healthcare providers cannot avoid.

The MPE's role also encompasses the ethical and data-governance dimensions of the regulatory framework. As noted in [Section 3.2](#), the radiology subspecialty curriculum already references the European Society of Radiology's recommendations on AI Act implementation and the joint European–North American multisociety statement on the ethics of AI in radiology [[7,8](#)]; the focus group recommends that comparable content be made common across all three subspecialties at the next curriculum revision. Assessing a clinical AI system for bias across patient subgroups, ensuring that training and validation data are representative of the local population, and contributing to the transparency obligations owed to patients and clinicians are all tasks in which the MPE's quantitative competency is directly applicable, and which correspond to Domain 4 of the advanced competency layer.

Two caveats frame this role. First, the MPE's regulatory contribution under the AI Act is necessarily collaborative rather than exclusive: responsibility for a high-risk clinical AI system is shared with the treating clinician, the manufacturer, the institution's data protection function, and, where applicable, the notified body — the MPE's role is to provide the technical assurance and oversight infrastructure on which these other parties rely, not to assume sole accountability for the system. Second, the precise national implementation of the AI Act's healthcare provisions, and the extent to which competent authorities formally recognize the MPE's role within it, remains a matter of ongoing development across EU member states at the time of writing. The focus group's recommendations in [Section 7](#) are framed with this evolving landscape in mind and are intended to position the profession to engage constructively with national implementation processes as they mature, rather than to describe a settled state of affairs.

7. Implementation recommendations

The preceding sections of this paper — the harmonization analysis ([Section 3](#)), the advanced and expert competency layer ([Section 4](#)), the CPD framework ([Section 5](#)), and the regulatory role of the MPE ([Section 6](#)) — generate a set of concrete actions that must be taken by specific bodies if the framework is to be implemented in European medical

physics practice. [Section 7](#) translates these findings into targeted recommendations addressed to four audiences: EFOMP and its governance bodies; NMOs; training centers and university programs; and individual MPEs and their employers. The recommendations are intentionally specific: they name the responsible body, describe the action, and indicate a realistic timeframe.

7.1. Recommendations to EFOMP

The focus group recommends that EFOMP: (1) commission the next round of subspecialty curriculum revisions with an explicit mandate to address the divergences and gaps identified in [Section 3](#); (2) establish a minimum common AI specification — covering at least the entry-level foundation of [Section 3.1](#), the EU AI Act, and AI ethics — that all three curricula must meet; (3) develop a formal recognition mechanism for Advanced and Expert AI competency, structured around the four domains of [Section 4](#); (4) coordinate with the European School of Medical Physics Expertise (ESMPE) on integrating AI lifecycle management into MPE certification; and (5) establish a standing AI curriculum review mechanism, faster than the five-yearly cycle governing the curricula as a whole, given the pace of change documented in [Sections 1 and 4.1](#). Given the urgency shown by the survey evidence in [Section 4.1](#), the focus group recommends initiating (1)–(3) within twelve months of publication.

7.2. Recommendations to NMOs

The focus group recommends that NMOs: (1) map national training programs against the harmonized EFOMP AI competency framework of [Sections 3 and 4](#); (2) establish or endorse CPD pathways for AI competency maintenance, consistent with the recurrent, accredited model of [Section 5](#), rather than relying on ad hoc course availability; (3) liaise with national competent authorities on the MPE's role under the EU AI Act, translating the profession-level position of [Section 6](#) into each member state's regulatory context — a task that will vary given the differing pace of national AI Act implementation; and (4) disseminate the Advanced and Expert competency framework to training centers and hospital departments ahead of any formal EFOMP recognition mechanism.

7.3. Recommendations to training centers and university programs

The focus group recommends that training centers and universities: (1) incorporate the harmonized AI core ([Section 3.1](#)) and correct the curricular divergences of [Section 3.2](#) at institutional level, without waiting for the next formal curriculum revision where local design permits; (2) introduce teaching on generative AI, foundation models, large language models and agentic AI ([Section 3.3](#), gaps 2–3) at both trainee and CPD levels, since these technologies are entering practice faster than the curricula address them; (3) integrate EU AI Act obligations and AI ethics as content common to all three subspecialties rather than radiology-specific material; and (4) begin offering the Advanced competency modules of [Section 4](#) as elective or post-qualification options, building institutional experience ahead of formal recognition.

7.4. Recommendations to individual MPEs and their employers

To individual MPEs, the focus group recommends: self-assessment against the competency framework of [Sections 3 and 4](#) rather than assuming general preparedness; structured CPD planning using the [Section 5](#) lifecycle framework, treating AI CPD as a recurrent obligation; and documentation of one's AI governance role under the EU AI Act within institutional deployment processes. To employers, the focus group recommends a clear institutional responsibility to provide

protected time and resources for AI training. This follows directly from the survey finding ([Section 4.1](#)) that managerial support for AI training is currently reactive — rising only once an institution has already adopted AI tools [[14](#)] — rather than anticipatory. The focus group considers this a risk to safe deployment and recommends that institutional support precede, or at minimum accompany, AI adoption rather than follow it.

7.5. From 2021 to 2026: what has changed, what has been superseded, and what is new

This subsection maps the three shifts described in [Section 1](#) — in AI technology, regulation, and the EFOMP curricular landscape — onto the specific content of the 2021 white paper [[1](#)], giving the focus group and the broader EFOMP community a clear audit trail between the two documents.

Aligned with the 2021 recommendation. The 2021 paper's central recommendation — that AI be integrated into the subspecialty curricula rather than treated as a permanent addendum — is aligned with the current curricula: all three now contain a dedicated core section on AI ([Section 3](#)). The published curricula do not document a direct causal link to the 2021 paper — no curriculum body's revision materials cite it or describe its influence on the process — so this alignment is presented as consistent with, rather than demonstrably caused by, the earlier recommendation; it may equally reflect the broader trajectory of AI adoption in healthcare that the 2021 paper itself anticipated. This is the clearest respect in which this paper's premise differs from 2021's: the entry-level gap the earlier paper addressed no longer exists in the form it did.

Updated or corrected. The 2021 paper's two-level Basic/Advanced structure is replaced here by a different layering rationale: a harmonized entry-level foundation delivered within the subspecialty curricula ([Section 3.1](#)), and a separate, subspecialty-agnostic advanced layer ([Section 4](#)) distributed across the professional lifecycle ([Section 5](#)) rather than fixed to one training point. The focus group judges this better reflects how AI competency is actually built and maintained, and aligns more naturally with EFOMP's existing CPD and expert-recognition structures.

New in this paper. Generative AI, foundation models, LLMs and agentic AI ([Section 4.3](#), Domain 3); the EU AI Act and the MPE's role under it ([Section 6](#)), which did not exist as law in 2021; a formal proposal for AI CPD accreditation ([Section 5](#)); AI lifecycle management as a distinct domain ([Section 4.3](#), Domain 2); and AI-specific cybersecurity ([Section 4.3](#), Domain 1) — none were addressed, or addressed in comparable depth, in 2021. Together these reflect the shift described in [Section 1](#): from establishing that AI matters to assuring it is used safely, effectively, and in regulatory compliance.

Carried forward unchanged. The 2021 paper's underlying position — that the MPE has a natural, necessary role in the safe deployment of AI, grounded in the profession's existing QA and regulatory competencies — remains fully valid and is carried forward as this paper's foundational premise. What has changed is not the premise but the technological, regulatory and curricular context in which it must now be enacted.

8. Discussion and conclusions

This paper set out to answer the question posed in [Section 1](#): what is the right framework for AI competency in medical physics now that the entry-level gap the 2021 white paper addressed has been substantially closed by the subspecialty curricula themselves? The answer has four parts: a harmonization analysis showing convergence on a shared methodological foundation alongside specific divergences and six common gaps ([Section 3](#)); a subspecialty-agnostic advanced competency

layer, grounded in five years of survey evidence, organized into four domains mapped to those gaps (Section 4); a professional-lifecycle model for acquiring, transitioning into, maintaining and formally recognizing this competency (Section 5); and an account of the MPE's increasingly formalized role under the EU AI Act, translated into recommendations for the four bodies best placed to act (Sections 6–7).

Together, these four elements mark a necessary departure from the 2021 framework. Where the 2021 paper argued that the profession needed to engage with AI, this paper starts from the premise that it does, and asks how that engagement should be structured, assured and regulated — a shift that reflects the field's maturation and the curricula's alignment with the earlier paper's recommendations (Section 7.5).

8.1. Limitations

The focus group notes five limitations. The harmonization analysis (Section 3) is based on the curricula as published, not on how they are actually implemented and assessed across European training programs, which may vary considerably. The comparative curriculum mapping and the framework's consensus basis (Section 2) rest on iterative expert deliberation and cross-checking within the Working Group, rather than on a named formal consensus method (e.g., a Delphi process) or an independently validated multi-coder protocol, and should be read as expert synthesis rather than as the output of a systematic, inter-rater-validated methodology. The survey evidence underpinning Section 4 draws on three surveys with partially overlapping populations and methodologies, and the most recent [15] was still being formalized for publication at the time of writing; findings from it should be treated as provisional. The four-domain structure and lifecycle model (Sections 4–5) are the focus group's synthesis and recommendation, not yet piloted or validated against outcomes — EFOMP should treat initial implementation as an opportunity for evaluation rather than a finished specification. Finally, this paper reflects the regulatory landscape at the time of writing; national implementation of the EU AI Act's healthcare provisions is still evolving, and the MPE's role described here may be clarified, in either direction, as that implementation matures. Some technologies discussed as advanced competencies may become standard components of entry-level training before the next curriculum revision.

8.2. Horizon scanning

AI capability has advanced faster than the profession's structures for assuring, governing and teaching it throughout the five years between the 2021 paper and this one. There is little reason to expect that pattern to reverse. Five developments seem to the focus group reasonably foreseeable, offered not as predictions but as a horizon against which this framework's adaptability — not just its current fit — should be judged.

First, generative AI, foundation models and agentic AI will likely continue moving from the periphery to the center of clinical and operational workflows, extending beyond today's image reconstruction and synthesis into report generation, decision support, and increasingly autonomous treatment planning. Agentic AI systems are expected to extend beyond content generation by autonomously orchestrating complex workflows through planning, tool use, API interactions and multi-step reasoning. In medical imaging, these systems are beginning to support AI quality assurance, predictive model development, deployment, continuous monitoring and lifecycle management, creating new competency requirements for MPEs in validating and governing autonomous AI pipelines. A brief mention of agentic AI has already been introduced earlier in this paper (Section 3.3); the increasing autonomy such systems carry is a development that will need sustained attention as it matures.

A second development worth watching is federated learning, an approach in which models are trained collaboratively across institutions

without any single party accessing another's raw data. Federated learning addresses long-standing data-sharing and privacy limitations that have constrained multi-institutional AI development in medical imaging, but it introduces its own competency demands: as Hatherley et al. [17] argue, federating a model across institutions creates a distinct form of opacity in which no single stakeholder can fully access, audit or curate the pooled training data, which complicates both accountability and security — federated models are, for example, more vulnerable to data-poisoning attacks than centrally trained ones. Validating and governing federated AI pipelines is therefore likely to become a related, but distinct, competency area for the MPE alongside the agentic AI developments discussed above.

A related, earlier-stage development worth naming is the emergence of digital twins and digital human twins in healthcare — patient-specific computational models built from imaging and biokinetic data that support individualized simulation of disease progression and treatment response [18]; should this mature into clinical use, its validation and governance would fall naturally within Domain 3 (Section 4.3), which is deliberately framed to extend to systems that do not yet exist and which the focus group expects, more than any other domain, to require the standing review mechanism of Section 7.1.

Fourth, the EU AI Act is a first-generation instrument for a fast-moving technology, and its healthcare application is still being worked out across member states (Section 6). The focus group expects further guidance, amendment and case law, and that the MPE's role under the Act will become more, not less, formally specified as practice accumulates — which is precisely why Sections 5 and 7 aim to position the profession to respond from demonstrated competency rather than be assigned an unprepared role.

Fifth, the entry-level/advanced distinction this paper draws may itself prove temporary. Just as generative AI and the AI Act were not writable into a 2021 curriculum, some of what this paper treats as advanced — AI lifecycle management especially — may be as foundational as basic statistics by the next curriculum revision. The focus group recommends treating that boundary as one to be reassessed at each standing review, not as a fixed architectural feature.

8.3. Conclusions

Over the past five years, artificial intelligence has evolved from an emerging technology to an integral component of modern medical physics practice. The profession's challenge has therefore shifted from introducing AI into specialist education to ensuring that AI competency is maintained throughout a professional career and translated into safe, trustworthy, effective and regulatory-compliant clinical practice. The framework proposed in this paper provides a foundation for that transition, but it is not intended to be static. As AI technologies, regulatory requirements and clinical applications continue to evolve, so too must the competencies of the Medical Physics Expert. The focus group therefore views this framework not as a final destination, but as the beginning of a continuously evolving model for AI competency within the profession. Ultimately, the objective of this framework is not simply to strengthen AI competency, but to ensure that the rapid evolution of AI continues to translate into safe, effective and high-quality care for patients.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper. Additional disclosure: Co-author Federica Zanca is affiliated with the European Innovation Council and SMEs Executive Agency (EISMEA). The views and opinions expressed in this manuscript are those of the authors alone and do not necessarily reflect the official policy or position of EISMEA or of the European Commission.

Appendix A. Comparative mapping of the three EFOMP subspecialty curricula

Feature	Radiology	Radiotherapy	Nuclear medicine
Curriculum section (AI/data content)	III.6	III.9	III.12
ECTS allocated to AI/data content	10	8	15
Shared statistical foundation (Section 3.1)	Yes	Yes	Yes
Shared core ML model families (Section 3.1)	Yes	Yes	Yes
“Perform QA of AI/ML models” as expert competency	Yes	Yes	Yes
Explicit coverage of transformers / foundation models / GANs / RNNs	No (CNN and classical ML only)	No (CNN and classical ML only)	Yes
Explicit treatment of the EU AI Act	Yes (cites [7])	No	No
Explicit treatment of AI-specific ethics	Yes (cites [8])	No	No
Reference to GDPR / MDR	Yes	Yes	Yes
Principal subspecialty-specific AI applications named	Scan planning, acquisition parameter optimization, image reconstruction and post-processing, segmentation, lesion detection, dosimetry	Pre-treatment QA, CBCT-to-Hounsfield-unit conversion, pseudo-CT generation from MRI, treatment planning, contouring, outcome prediction	Image reconstruction, tumor volume segmentation, pseudo-CT generation for PET/MR attenuation correction, low-dose imaging, predictive modeling, radiomics

Note: this table summarizes the analysis presented in Section 3.1–3.2. Convergent rows (shared statistical foundation, shared ML model families, QA competency, GDPR/MDR reference) represent the harmonized core described in Section 3.1; divergent rows (ECTS, architecture coverage, AI Act/ethics treatment) are discussed in Section 3.2 together with the focus group’s recommendations for the next curriculum revision.

Appendix B. Mapping of the six curricular gaps to the four advanced competency domains

Gap (Section 3.3)	Advanced domain (Section 4.3)	Addressed in
1. AI lifecycle management beyond commissioning	Domain 2: AI lifecycle management	Section 4.3; Section 6
2. Generative AI and foundation models, including LLMs	Domain 3: Generative AI and foundation models	Section 4.3
3. Agentic AI systems	Domain 3: Generative AI and foundation models	Section 4.3
4. Operationalizing human oversight under the EU AI Act	Domain 4: AI governance	Section 4.3; Section 6
5. Post-market surveillance methodology specific to AI	Domain 1: Advanced validation and QA; Domain 2: AI lifecycle management	Section 4.3; Section 6
6. AI-related cybersecurity	Domain 1: Advanced validation and QA	Section 4.3

Appendix C. Mapping of survey training priorities to competency domains

The eight training topics ranked by respondents to the 2025 EFOMP AI SIG survey (Section 4.2), shown in descending order of interest, map to the advanced competency domains and to the co-selection clusters identified in the survey’s Jaccard similarity analysis.

#	Training topic (ranked)	Advanced domain	Cluster
1	Validation of AI performance	Domain 1: Advanced validation and QA	Clusters 1 & 3
2	Medical image analysis with AI	Domain 1 (applied)	Cluster 1
3	The role of the medical physicist in AI	Cross-cutting (varies significantly by country)	Cluster 1
4	Big data analysis	Entry-level core (Section 3.1)	—
5	Basics of machine and deep learning	Entry-level core (Section 3.1)	—
6	Regulatory issues in AI	Domain 4: AI governance	Clusters 2 & 3
7	Advanced machine and deep learning	Domain 3: Generative AI and foundation models	—
8	Ethical issues in AI	Domain 4: AI governance	Cluster 2

Cluster 1: validation × image analysis × MPE role (applied assurance). Cluster 2: regulatory × ethical (governance). Cluster 3: regulatory × validation (bridging). Only “the role of the medical physicist in AI” showed statistically significant variation by country [15]. See Section 4.2.

References

[1] Zanca F, Hernandez-Giron I, Avanzo M, Guidi G, Crijns W, Diaz O, et al. Expanding the medical physicist curricular and professional program to include artificial intelligence. *Phys Med* 2021;83:174–83. <https://doi.org/10.1016/j.ejmp.2021.01.069>.

[2] Kim C, Gadgil SU, DeGrave AJ, Omiye JF, Cai ZR, Daneshjou R, et al. Transparent medical image AI via an image–text foundation model grounded in medical literature. *Nat Med*. 2024;30(4):1154–65.

[3] Wang J, Wang K, Yu Y, et al. Self-improving generative foundation model for synthetic medical image generation and clinical applications. *Nat Med*. 2025;31(2):609–17.

- [4] EFOMP core curriculum for medical physics experts in radiology, 3rd Draft (in press, 2026).
- [5] EFOMP core curriculum for medical physics experts in radiotherapy, 3rd Revised Edition, 2022.
- [6] EFOMP, core curriculum for medical physics experts in nuclear medicine, 2nd revised edition, 2024.
- [7] Kotter E, D'Antonoli TA, Cuocolo R, Hierath M, Huisman M, Klontzas ME, et al. Guiding AI in radiology: ESR's recommendations for effective implementation of the European AI Act. *Insights. Imaging* 2025;16:33. <https://doi.org/10.1186/s13244-025-01905-x>.
- [8] Geis JR, Brady AP, Wu CC, Spencer J, Ranschaert E, Jaremko JL, et al. Ethics of artificial intelligence in radiology: summary of the joint European and North American multisociety statement. *Insights Imaging* 2019;10:101. <https://doi.org/10.1186/s13244-019-0785-8>.
- [9] Hurkmans C, Bibault JE, Brock KK, van Elmpt W, Feng M, David Fuller C, et al. A joint ESTRO and AAPM guideline for development, clinical validation and reporting of artificial intelligence models in radiation therapy. *Radiother Oncol.* 2024 Aug;197:110345.
- [10] Ferber D, Hilgers L, Höper C, et al. Towards autonomous medical artificial intelligence agents. *Nature* 2026. <https://doi.org/10.1038/s41586-026-10675-5>.
- [11] Flanders AE, Peng Y, Prevedello L, et al. A multi-agent large language model framework to automatically assess performance of a clinical AI Triage tool. *npj Health Syst* 2026;3:35. <https://doi.org/10.1038/s44401-026-00100-4>.
- [12] Tzanis E, Klontzas ME. ReclAlm: a multiagent framework for monitoring and correcting performance decline in medical imaging AI. *Radiol Artif Intell.* 2026;8 (4):e250923. <https://doi.org/10.1148/ryai.250923>.
- [13] Diaz O, Guidi G, Ivashchenko O, Colgan N, Zanca F. Artificial intelligence in the medical physics community: an international survey. *Phys Med* 2021;81:141–6. <https://doi.org/10.1016/j.ejmp.2020.11.037>.
- [14] Doolan P, Michopoulou S, Meades R. IPEM topical report: results of a 2024 UK survey of artificial intelligence in medical physics and clinical engineering. *Phys Med Biol* 2025;70(14). <https://doi.org/10.1088/1361-6560/ade92b>.
- [15] EFOMP AI Special Interest Group. Artificial intelligence in medical physics: results of a 2025 European survey of the EFOMP AI Special Interest Group. Manuscript in preparation, 2025.
- [16] European Commission. Regulation (EU) 2024/1689 of the European parliament and of the council laying down harmonized rules on artificial intelligence (artificial intelligence act). Official Journal of the European Union, 2024.
- [17] Hatherley J, Søgaard A, Ballantyne A, Pauwels R. Federation opacity and the promise of federated learning in healthcare. *Am J Bioethics* 2026;1–14. <https://doi.org/10.1080/15265161.2026.2637093>.
- [18] Papachristou K, Katsakiori P, Papadimitroulas P, Strigari L, et al. Digital twins' advancements and applications in healthcare, towards precision medicine. *J Pers Med* 2024;14(11):1101. <https://doi.org/10.3390/jpm14111101>.