The 'Wild West' of Medicine: Exploring the Emergence of 'Grassroots' AI Governance in Radiology

Bhargavi Ganesh[1,2,3], Daniel S. Schiff[3], and Stuart Anderson[2]

[1]Centre for Technomoral Futures, University of Edinburgh

[2]Department of Informatics, University of Edinburgh

[3]Governance and Responsible AI Lab, Purdue University

b.ganesh@sms.ed.ac.uk, dschiff@purdue.edu, s.anderson@ed.ac.uk

Abstract

In the past 10 years, a steady increase in clinical AI adoption has been accompanied by concerns regarding potential risks. Thus, there has been a growing body of literature on the regulatory implications of AI devices, and studies exploring clinician attitudes towards AI. However, there has been limited work examining 'bottom-up' hospital-level AI governance approaches. To fill this gap, we conducted a qualitative study interviewing 22 healthcare practitioners with AI governance experience within radiology departments and/or professional societies in the US and UK. We aimed to understand the current state of AI adoption and governance, clinician perspectives on responsibility, and the interaction between 'top-down' and 'bottom-up' governance approaches. Our findings indicate disparities in resources and AI expertise, as well as differences in the scope, composition, remit, and role of AI governance committees across hospitals. Additionally, we uncover emerging challenges in negotiating responsibility norms for AI outcomes and performance monitoring. We also discuss the AI governance role taken on by some clinicians, often on a voluntary basis, and the challenges they face in navigating siloed, hierarchical organizations. Finally, we analyze participant recommendations, including the development of streamlined guidance on responsible AI adoption, better staff education/training, and centralized approaches to performance monitoring.

Introduction

In the past 10 years, AI has rapidly grown in the healthcare sector. In the US, for example, over 100 new AI-driven medical devices have been approved by the Food and Drug Administration (FDA) every year since 2020. Although techniques for automating image processing and feature identification have been developed and used since the 1990s, this recent growth in the development of AI-driven devices has been driven by developments in neural networks and transformer models, the availability of more compute resources, and increased adoption of electronic health records globally.

AI-driven medical devices have been developed for a number of different clinical purposes, including earlier detection of abnormalities like tumors, improving efficiency and clinician confidence by acting as a "second pair of eyes", better management of staff resources by triaging and prioritizing the most urgent cases, and even improving the speed and quality of image generation. With the increased accessibility and improved performance of large language models (LLMs), the use of AI has also expanded to administrative tasks like patient scheduling (Knight et al. 2023), summarizing clinical notes (Van Veen et al. 2024), and patient messaging (Chen et al. 2024; Garcia et al. 2024). Despite the positive potential of these systems, however, there have been concerns regarding automation bias (Abdelwanis et al. 2024; Dratsch et al. 2023), gender or racial biases (Gichoya et al. 2022; Adam et al. 2022), patient privacy (Murdoch 2021), worsened clinician burnout (Kwee and Kwee 2021), and other risks.

Preceding the adoption of newer 'AI' systems, clinical decision support systems (CDSS) have been widely used in hospitals. There are a number of studies on clinician attitudes towards CDSSs (Pope et al. 2013), implementation challenges they generate (Liberati et al. 2017; Ford et al. 2021), and their impact on patient outcomes (Garg et al. 2005). Recently, there have also been an increasing number of qualitative studies exploring clinician attitudes towards AI, including perceived benefits and concerns (Oremule et al. 2024; Fazakarley et al. 2023; Nittas et al. 2025; Parikh et al. 2022), and perspectives on how responsibility should be distributed for AI outcomes (Laï, Brian, and Mamzer 2020; Van Cauwenberge et al. 2022).

Scholars have also discussed regulatory (top-down) governance challenges generated by AI in healthcare (Palaniappan, Lin, and Vogel 2024; Ong et al. 2025; Gilbert

et al. 2023; Castonguay et al. 2024). However, with the exception of a few recent studies (Nong, Hamasha, and Platt 2024; Chow, Lee, and Wu 2025) there is limited research on how AI is currently being deployed and governed at individual hospitals. This 'bottom-up' aspect of governance is particularly interesting in the context of healthcare, as the highly regulated nature of the sector has created complex dynamics between stakeholders within different layers of governance, including hospitals, professional societies, and regulators at the regional or national level. In our paper, we build on an 'institutional approach' to healthcare governance (Abimbola et al. 2017), to explore the relationship between bottom-up and top-down approaches to AI governance.

To that end, we conducted qualitative interviews with 22 participants involved in AI governance within radiology departments in the US and UK; focusing on this subspecialty and these countries as they have the greatest maturity of AI adoption (Castonguay et al. 2024; Allen et al. 2021; US Food and Drug Administration 2025). We conducted semi-structured qualitative interviews with radiologists, radiation oncologists (also known as medical physicists and clinical scientists), and researchers based in radiology departments; prioritising individuals with experience on AI governance committees within their hospital and/or within a medical professional society. We also interviewed a small sample of industry participants and regulators with experience in AI development and governance. Across these interviews we aimed to answer the following research questions: (**RQ1**) What is the current state of adoption and governance of AI in radiology departments in the US and UK? (**RQ2**) How do clinical staff in radiology perceive their role/responsibility for AI-related outcomes? (**RQ3**) Does AI pose unique challenges at the intersection between bottom-up governance and top-down rules? (**RQ4**) What recommendations do experts have for improving governance of AI going forward?

Using abductive thematic analysis, we uncover the following themes. First, we find variance in the types of governance mechanisms designed at different hospitals, owing to different perceptions of AI, availability of resources, and technical expertise. This finding suggests that the specialty-driven approach in medicine is increasingly at odds with the deployment of AI across a wide range of contexts. Second, our findings show that some clinicians have taken on additional AI-focused governance roles, often on a voluntary basis, either within their own hospitals or as part of task groups and working groups formed within professional bodies. We discuss these roles within the context of 'policy entrepreneurs' (Durant and Diehl 1989), while noting the unique challenges that participants face in operating within deeply hierarchical organizations, navigating power relations, and justifying their role. Third, we find that although participants almost unanimously agreed that the clinician using an AI device bears ultimate responsibility for patient outcomes, there was ambivalence regarding who should be responsible for ongoing performance evaluation of AI devices; suggesting the practical risk of clinicians being placed in 'moral crumple zones' (Elish 2019). Finally, our results demonstrate participants' simultaneous desire for standardized guidance on how to adopt AI, and robust local governance measures tailored to unique patient populations; indicating possible tensions between top-down and bottom-up approaches towards AI governance.

Background and Theoretical Underpinning

Prior to interviewing participants, we created a topic guide, developing interview questions from theoretical frameworks on responsibility and institutional governance. Below, we provide some detail about these frameworks.

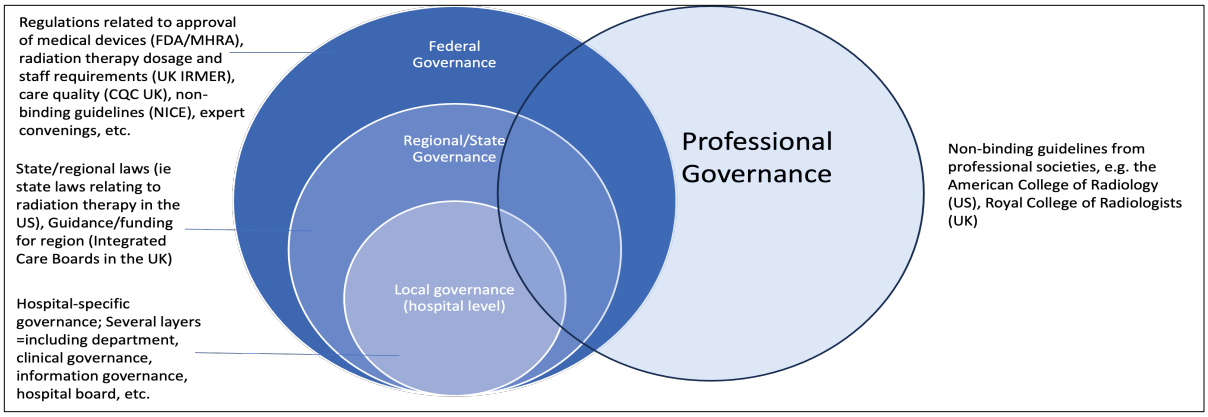

Figure 1: Overlapping layers of healthcare governance

Novel Responsibility Challenges Generated by AI

Ever since Andreas Matthias coined the term "responsibility gap," (Matthias 2004) scholars at the intersection of AI ethics, policy, and governance have been using this concept to articulate the challenge of assigning responsibility for the outcomes of AI systems (Coeckelbergh 2020; Nyholm 2018; Santoni de Sio and Mecacci 2021; Bleher and Braun 2022). In particular, the argument is that since machine learning algorithms automate some aspects of decision-making, it is challenging to assign responsibility to the human operators of such systems. This challenge can theoretically create a 'gap', where no one is ultimately responsible for AI outcomes despite society continuing to pay the cost for its harmful impacts (Matthias 2004). On the other hand, some scholars highlight the risk that ambiguity in responsibility attribution can lead to one person absorbing all the responsibility for an AI outcome, thus falling within the 'moral crumple zone' (Elish 2019; Smith and Fotheringham 2020). Scholars also note that the increasingly complex supply chains involved with AI development and deployment (Widder and Nafus 2023; Cobbe, Veale, and Singh 2023) generate challenges for identifying which of the 'many hands' (Coeckelbergh 2020; Nissenbaum 1996) should be held accountable for AI-related outcomes.

Within the context of healthcare, although there has been some theoretical work suggesting that AI systems used for clinical decision-making generate novel challenges for assigning responsibility (Bleher and Braun 2022), empirical studies have shown the continued emphasis clinicians place on having the 'final responsibility' over clinical outcomes when using AI for decision support (Jones, Thornton, and Wyatt 2023; Van Cauwenberge et al. 2022; Laï, Brian, and Mamzer 2020). Thus, in our qualitative study, we aimed to understand participants' perceptions of their own responsibility for both AI use and governance within the hospital.

Relationship Between 'Bottom-Up' and 'Top-Down' Governance

In our analysis of AI governance, we refer to both formal regulations and legal liabilities, and informal guidelines such as ethical codes of conduct, practice guidelines, and informal governance committees. Governance has historically been one means of addressing 'gaps' in responsibility emerging from the deployment of new innovations (Vallor and Ganesh 2023). This can be done through forward-looking assignation of responsibility, such as assignment of duties, roles and offices to morally and legally answerable individuals, or backward-looking inspection of how causal, moral and legal responsibilities should be assigned. Because healthcare is a heavily regulated field, with binding legal and regulatory obligations, as well as an established culture of responsibility and ethical norms at the individual, organizational, and professional level (Hajar 2017; Bryden and Storey 2011; Nandi 2000; Beauchamp and Childress 1979), there is already a precedent for robust forward and backward-looking assignation of responsibility, with clearly delineated roles, hierarchies, and protocols for quality assurance within hospitals (Jalilvand, Raeisi, and Shaarbafchizadeh 2024; Liberatore and Nydick 2008; Esposito and Dal Canton 2014; Foy et al. 2020).

However, the different layers of healthcare governance create complex dynamics for stakeholders governing AI use. On one hand, there are regulatory bodies, such as the FDA, and the Medicines and Healthcare products Regulatory Agency (MHRA) in the UK, which set requirements for medical device testing before they are placed on the market. In addition, some regions have specific laws, such as regulations on the use of X-ray equipment within that region. Within hospitals, hospital boards and administrators are legally responsible for monitoring quality and safety, but delegate this authority to clinical committees or staff councils (Jalilvand, Raeisi, and Shaarbafchizadeh 2024). Governance also emerges at the level of professional societies, which are often subspecialty-specific, and develop targeted practice guidelines for adoption by individual hospital departments. A simplified diagram of these different relationships is presented below in Figure 1.

AI generates some novel governance challenges when compared to previous technologies. For one, because the performance of AI systems is so dependent on the data used to train a given algorithm, performance can vary across different patient populations (Sahiner et al. 2023). Thus, regulatory approval does not guarantee that an AI system will perform well on a given hospital's data. Currently, regulatory approvals are also restricted to 'static' systems that do not evolve over time, generating risks that AI system performance degrades over time (Sahiner et al. 2023). In response to these risks, in addition to subjecting AI medical devices to thorough testing requirements prior to placing them on the market, regulatory bodies like the FDA and the MHRA have also begun developing guidance for reporting post-market performance, known as post-market surveillance. Despite these interventions, however, regulatory bodies still have limited oversight over the use of AI in clinical practice. For example, long regulatory approval times have incentivized companies to apply for approval in lower-risk categories, assuming that clinicians will use these systems differently 'off-label' (Krishnamoorthy, Sjoding, and Wiens 2024). To fill these gaps, there have been an overwhelming number of professional societies, standards bodies, and other related organizations issuing guidelines on the clinical adoption of AI (Brady et al. 2024; Hadjiiski et al. 2023; The Royal College of Radiologists 2024; British Medical Association 2024; American Medical Association 2024). Some individual hospitals are also now developing processes and guidelines for adopting AI (Kim et al. 2023; Daye et al. 2022).

Although our qualitative study is primarily concerned with bottom-up governance, the intersection between various layers of governance motivated us to explore an 'institutional approach' to governance, focusing on the relationship between hard rules and emerging 'building blocks' at the local level (Abimbola et al. 2017). To that end, our interview topic guide included questions related to local, professional, and regulatory governance. We discuss the interview methodology in more detail below.

Interview and Coding Methodology

The expert interview is commonly used in the social sciences for exploratory analysis, obtaining of specialized knowledge, or theory generation (Bogner, Littig, and Menz 2009). For our study, semi-structured interviews presented advantages over other methods, such as surveys, because they enabled a more in-depth understanding of individual motivations and organizational dynamics. Additionally, given the sensitive regulatory and legal landscape of the healthcare sector, and the relatively early stage of AI adoption and governance, information and rationales for hospital-specific AI governance measures would be difficult to assess without consulting with experts who are involved with implementing these measures on the ground.

Recruitment Strategy and Participant Summary

Participant recruitment focused on 'experts', who defined as individuals with experience using or governing AI in radiology departments, including radiologists, radiation oncologists (also known as medical physicists and clinical scientists), and researchers. The field of radiology was chosen as it is currently the most mature area of AI adoption in healthcare (Allen et al. 2021). We interviewed participants in the US and UK, as these countries currently lead in AI development and adoption among OECD countries (Castonguay et al. 2024). In addition to clinical participants, we interviewed a subset of participants from regulatory and industry backgrounds, to explore the interaction between local, professional, and regulatory governance. Participants were recruited through professional contacts, radiology and radiation oncology conferences, advertising on social media, and snowball sampling (Vogt and Johnson 2015). All of the clinicians (including those not involved with governance) had experience using AI in their day-to-day work. Exactly half of the participants were US-based and the other half were UK-based. Roughly 2/3 were male, while 1/3 were female (mirroring gender imbalances in radiology). The majority of US participants (all but one) belonged to academic medical centers, which are typically 'early adopters' of new technologies (Faruki, Zane, and Wiler 2022), and only three UK participants (representing 2 hospital systems) were practicing at 'Academic Health Science Centers' (AHSCs).

Role	
Radiologist	8
Radiation Oncologist	8
Clinical Researcher	3
Industry	2
Former Regulator	1

Table 1: Summary of Participant Role/Professional Title

Governance Involvement	
Local	5
Professional	5
Local and Professional	5
Regulatory	1
None	6

Table 2: Summary of Participant Governance Activity

Interview Structure and Topic Guide

Interviews were conducted by the primary author from March 2024 – March 2025, for 45-60 minutes (each) on Microsoft Teams. Two pilot interviews were conducted with clinicians (one each from the US and UK) experienced with using AI in their day-to-day practice. Questions were derived from the theoretical frameworks on responsibility and institutional governance presented above. The topic guide was refined after pilot interviews and peer debriefing (Lincoln, Guba, and Pilotta 1985) with co-authors. Participants were asked questions about their background/training, experiences and benefits/ concerns with using AI, individual responsibility for AI outcomes, involvement in professional/local governance (if applicable), thoughts on regulatory governance, and general recommendations on how AI governance could be improved in radiology.

Coding and Analysis

We utilized abductive thematic analysis (Thompson 2022), to undertake an iterative process of triangulating between data and theoretical frameworks, also known as 'theoretical triangulation' (Morse, Knafl, and Breitmayer 1991; Denzin 2017). All coding was done using the qualitative coding software NVivo (14). The codebook was initially developed by the primary author and applied to two pilot interviews. The co-authors then reviewed the codebook and provided suggested revisions, which mainly involved simplifying and condensing the number of codes, and adding in demographic codes. The primary coder finished coding the re-

maining interviews, and went through a second round of reviews, in which one co-author randomly selected two interviews (one each from the US and UK) to review. On the basis of this review, the primary coder refined the codebook and methodology, after which they ensured that the methodology was uniformly applied across interviews, to establish 'internal consistency' (Guba 1981). The final codebook had 5 main code families, 36 sub-codes and 144 codes.

Results and Findings

AI as the 'Wild West' of Medicine

The first key finding was that strategies for adoption and governance of AI devices varied widely across the 15 institutions represented in our interview sample. This finding was surprisingly consistent across the US and UK, despite the UK's relatively more centralized NHS system.

UK participants noted that in order to incentivize greater adoption of AI, the government fully funded stroke detection AI devices within NHS trusts in England in 2020, and later set up an AI Diagnostic Fund (GOV.UK. 2023), which provides funding for NHS trusts to deploy a wide range of AI systems. Despite this centralized funding stream, however, participants suggested that the way that these devices were adopted and governed depended on the availability of individual hospital resources:

> "I think the current way we do it is you've got money coming down from above…to support projects and then…[it's] developed locally and you've got very different approaches to it depending on resources and need…I worry that…some places are doing it in a way that…is slightly dodgy because they just don't have…the infrastructure or the resources to do it well." (Radiologist, non-AHSC, UK).

Although US academic medical centers tended to be well-resourced, participants still highlighted variance in whether or not a hospital has a governance policy for AI.

> "Oh yeah, AI is the Wild West in medicine…So some hospitals may not have any governance at all. We're developing a governance policy, so technically we don't…have one…actively in use right now… same thing goes with other hospitals I've worked at." (Clinical Medical Physicist, Academic Medical Centre, US).

In addition to differences in resources, participants speculated about some of the other reasons for this variation across hospitals. One radiologist from a US academic medical center suggested that the 'slow boil' nature of AI adoption makes it appear deceptively low risk when compared with previous technologies:

> "AI…has been so insidious that… it's become integrated…piece by piece… it's a…slow boil…we didn't expect the radiologist…who's never been trained on a CT [to] say, oh yeah, I know how to use this… [with AI] practices [are]…caught unaware… you realize… we have 5 tools and we're not really tracking quality for any of these.. [and seeing] if this is being used effectively." (Radiologist, Academic Medical Center, US).

In line with this observation, some participants questioned whether AI should be viewed as a medical device in need of strict guardrails and monitoring, or just 'another piece of software' (Clinical Scientist, non-AHSC, UK), or 'software plugin' (Radiologist, Academic Medical Centre, US). Despite some variance in participant attitudes regarding the level of risk posed by AI systems, and the rigorousness of governance needed to mitigate these risks, participants commonly pointed out that AI use generated new concerns regarding overreliance, bias, performance on local patient data, lack of AI expertise among deployers/users, lack of explainability, and performance drift over time.

Responsibility 'Negotiation' Across Stakeholders

Across our interviews, we found that although participants commonly believed that the clinician has the 'final responsibility' over AI-related outcomes, in line with previous studies (Jones, Thornton, and Wyatt; Van Cauwenberge et al. 2022), AI deployments had also necessitated the negotiation of new responsibility norms within the hospital.

Participants typically named the responsible person as the one who had 'signed the report', 'acted on the report', or 'made a decision with the tool', while in some cases they named a specific clinical stakeholder like the radiologist or clinician. Less commonly, participants referred to hospitals as the main responsible party, or mentioned shared responsibility with the vendor, while one radiation oncologist stated that no one should be responsible, "if everyone's done their best and acts in good faith" (Clinical Scientist, non-AHSC, UK). Participants who viewed clinicians as having 'final responsibility' often cited legal liability, while others cited professional norms, duties to patients, or vendors' unwillingness to take on this responsibility. These perspectives suggest the emerging risk of 'moral crumple zones', particularly if clinicians take on responsibility for outcomes for which they lack either knowledge or control (Elish 2019; Smith and Fotheringham 2020).

Even those who were initially confident about clinicians having the 'final responsibility', later acknowledged that the answer would be less straightforward if AI was used more autonomously. A few US radiologists said that they would have to look to the automotive industry to understand how this would play out legally. In some cases, when an AI system was used for diagnosis by non-radiologists, new responsibilities could be added to existing roles:

> "We've had cases where the AI has said…there's a hemorrhage here. And it's been reviewed by a radiologist who's said that there isn't. But the clinician sees the AI results… [and] either acts on that…or they'll question the radiologist's opinion…I heard of a case [where] the physician insisted that the radiologist call a neuroradiologist at a different center. And this is in the middle of the night to confirm that there was no hemorrhage and so it led to a whole load of extra work" (Radiologist, AHSC, UK).

Participants also expressed ambivalence about who should be responsible for evaluating AI system performance over time; a finding that was consistent with a recent study (Chow, Lee, and Wu 2025).

> "So…there's a couple of different routes we could take [for monitoring]…but…it's kind of no point in going too far down the road until someone agrees whose responsibility it is...Because if you have hundreds of AI tools…that's a big job...coming into the trust. Should it be a central group? Should it be…each individual department that takes responsibility? I don't know the answer really." (Clinical Scientist, non-AHSC, UK)

Some participants asserted that it was the responsibility of individual departments to monitor performance, while others suggested that there should be a more centralized body, such as the NHS, or an independent third party, in charge of performance monitoring. A US radiologist mentioned that it was the responsibility of their hospital's AI governance group to work with different sections to put together viable test data prior to deploying the system to a small group of test users (Radiologist, US, Academic Medical Center). Another US radiologist stated that due to a lack of resources, their approach was to trust vendors to report performance, but added that ideally it would be better to have an independent body conducting evaluations (Radiologist, US, non-Academic Medical Center). These perspectives suggest that performance evaluation is approached very differently based on resources and capacity, and point to potential 'gaps' that emerge if no one is ultimately taking on this responsibility.

Designing 'Bottom-Up' Governance Approaches

Across our interviews, we observed that hospital-level governance committees tended to be formed through a 'bottom-up' process driven by motivated clinical staff (who we refer to as governance entrepreneurs). These individuals often wrestled with key questions regarding the appropriate scope, composition, remit, and role for their committee, and highlighted the importance of counteracting perceived inadequacies in stakeholder engagement, procurement processes, pre-deployment testing, and/or ongoing performance monitoring. Participants who were part of working groups developing best practice guidelines within professional societies tended to describe a similar "bottom-up" approach.

While some governance roles were formally included as part of participant job descriptions, many governance functions were undertaken on a completely voluntary basis, sometimes by temporary employees such as medical students and residents. Thus, much like regulatory "policy entrepreneurs" or "political actors who engage in collaborative efforts in and around government to promote policy innovations" (Durant and Diehl 1989), these governance entrepreneurs exhibit a similar "willingness to invest their resources – time, energy, reputation, and sometimes money – in the hope of a future return" (Durant and Diehl 1989). Governance entrepreneurs in our study referenced "goodwill," or doing the work "because we think it's the right thing to do, not because there's money attached or any resource" (Clinical Scientist, UK, non-AHSC).

Many of the participants pushing for more formalized AI governance processes within their radiology department were radiation oncologists with scientific backgrounds. For these participants, governance appeared to represent a chance to have a seat at the table, and provide expertise. As one participant noted:

> "You don't tend to get our profession at senior levels and trusts, so they wouldn't even know to come and find us... lots of people I know do complain about being essentially overlooked…when they could help." (Clinical Scientist, UK, non-AHSC)

One challenge faced by governance entrepreneurs was determining the scope of governance, given the use of AI for a wide range of different tasks:

> "One thing that concerns me is not being general enough...I've noticed people tend to fixate on things that…impact their day-to-day profession. So like radiologists focus more on CAD AI systems. Physicists might be more focused on…image acquisition...I've identified that there's these sort of niche groups. My bigger concern is policies that can…govern all AI in…radiology." (Medical Physicist, Academic Medical Centre, US)

Another participant viewed the increasing scope as a natural response to the more widespread use of AI:

> "[Our committee] started out as AI in radiology and we've transitioned that more to AI in imaging because all of our tools are imaging based. But it is possible that you know that will eventually be rebranded again…eventually it'll be an AI in medicine, as we continue to change and bring on more applications." (Radiologist, Academic Medical Centre, US)

Participants also highlighted the challenge of deciding which stakeholders should be involved in their department's governance committee. One US radiation oncologist described a committee involving 'physicians, imaging scientists…nursing staff…representation from IS and IT…and an ethicist' (Radiation Oncologist, US, Academic Medical Centre). Another US radiologist mentioned the involvement of staff with expertise on the 'clinical side…quality assurance side…information security side…development side' (Radiologist, US, Academic Medical Center). This participant also mentioned that their governance committee was modeled off of recommendations made by a paper (Daye et al. 2022) published in a medical journal. A UK radiation oncologist mentioned the inclusion of patients as 'part of the [governance] panel, and also part of a separate parallel panel' (Radiation Oncologist, UK, AHSC), while a US radiation oncologist expressed interest in considering different stakeholders, including the section chief, and ethical/regulatory experts, but worried about how 'efficient the committee would work if we even introduced all these different people' (Medical Physicist, US, Academic Medical Center)

Governance entrepreneurs also deliberated about what information their committee should request from clinicians interested in purchasing a new AI tool. This concern also played a role in participants' views of their committee as either being a 'gatekeeper' or merely providing 'oversight'.

> "Our committee is mostly not in the business of…rejecting things typically. Usually the idea is to push back and say this you need to do a better job of documenting this piece or justifying that or thinking about…this... but I think there's…concern that if they're too…forcefully saying no …people are gonna find ways to go around it as opposed to…going through it" (Radiation Oncologist, Academic Medical Center, US)

Perspectives on the role of the committee differed across hospitals, however. The participant referenced above noted that committees at other hospitals may envision their role differently, with some considering the goal to be to 'reject things they don't think will be useful', and act as a 'toll booth.' Echoing this sentiment, one UK radiation oncologist explicitly referred to their committee process as a "gateway…primarily concerned with commercial tools," (Clinical Scientist, UK, non-AHSC) and multiple participants (from US Academic Medical Centers) mentioned several rounds of acceptability thresholds that needed to be met for a tool to be deployed.

Governance entrepreneurs also faced challenges integrating new frameworks within existing policies. In particular, UK-based practitioners discussed challenges with conducting rigorous due diligence of AI systems, due to existing information governance and procurement laws.

> "Procurement is very rigid. There's not really a mechanism for evaluating how well…each individual tool works [on] your own dataset locally. You just have to…ask them for evidence and score it based on that...we're trying to persuade our procurement colleagues…to do something different for AI…I think it's going to be hard…because…the system they use [is the]…legally agreed way to procure things." (Radiologist, UK, non-AHSC)

The desire to test AI systems on local data before deployment has led to the design of pilot programs or trial periods, but hospitals diverged in their approach to these as well. Trial periods could involve deploying the AI system in 'shadow mode', meaning that the AI system is working in the background and collecting data on performance metrics. In some hospitals, these trials were followed by small rollouts and feedback from focus groups. However, several participants mentioned manufacturers' reluctance to offer trial periods, while some mentioned challenges in determining how long trial periods should be.

Participants noted that merely testing the AI system before deployment does not guarantee stable performance over time. Thus, designing performance monitoring guidelines was also a common challenge expressed by participants. One participant highlighted the lack of a 'gold standard' for evaluation:

> "Only a handful of places could even attempt being able to really do a rigorous evaluation…and then the question of course becomes who's the gold standard? Maybe the experts didn't see anything, but maybe [there are] hidden things that the AI is using to make its determination." (Radiologist, Academic Medical Centre, US)

In addition to designing new guidelines, governance entrepreneurs also have to navigate challenges in enforcement, particularly when aspects of governance, such as performance monitoring, are devolved to individual departments.

> "I think the main risk we face clinically is that so much work goes into navigating the initial governance problems [that] post deployment surveillance is…an afterthought…they go ok so you want me to do what?...give you… monthly reports…demonstrating the ongoing efficacy or… validity of the model...Whereas…if these were things done by a staff actor, then of course…we would have audit, we would have review...appraisal mechanisms." (Radiation Oncologist, AHSC, UK)

Finally, some governance entrepreneurs highlighted the need to find funding for their operations. While larger institutions with resources and buy-in from higher ups did not face this problem, some participants in the UK mentioned applying for short-term grants, through AI-specific fellowships or charitable funding bodies.

Suggested Governance Actions

Streamlined Guidance for Responsible Adoption

Participants commonly mentioned that there were too many regulatory and professional bodies involved in setting guidelines for AI.

> "I mean, everybody's sort of like the dog in the neighborhood that pees all over and wants to claim as much…of the neighborhood as his domain as possible. So I think there are lots of people that have done lots of things with governance." (Radiologist, Academic Medical Center, US)

> "There's almost too many [governing bodies] involved, with competing or overlapping scope and priorities. But in terms of...NHS E, NHS digital or all the different…colleges and institutions and stuff…and MHRA and the government…" (Clinical Scientist, UK, non-AHSC)

One US radiation oncologist expressed concerns around the 'confusion' generated if people are unsure which governance guidelines to use (Medical Physicist, US, Academic Medical Center), while another UK radiologist called on the NHS to have one dedicated body focused on ensuring AI is deployed safely NHS-wide (Radiologist, UK, non-AHSC).

Education/Training Across Hospital Staff

Participants commonly noted the importance of better education and training for clinicians/staff procuring and using AI systems. They highlighted that education could ensure that clinicians are not misled by vendor marketing.

> "When I've been on a call with a vendor, they pushed heavily that their tool was CE marked… a radiologist who I spoke to…when I said… just because it's CE marked doesn't mean it works [on your hospital's data]. [They were] like, well, I thought that is what that meant. And [they've] deployed some AI tools in [their] trust." (Clinical Scientist, UK, non-AHSC)

Participants also suggested that education was needed to prevent users from blindly relying on AI systems. A US-based researcher likened the need for education to the need to pass a driving test (Researcher, US, AHSC), while a US radiologist suggested that AI should be treated as a core competency within radiology education, and included on board exams (Radiologist, US, AHSC). Participants differed in their assessment of who should be responsible for education, with some mentioning hospitals themselves, others mentioning professional societies, and still others suggesting the need for government involvement.

Knowledge Sharing

Participants often suggested that better knowledge sharing was needed across institutions.

> "There's really not a Consumer Reports that exists right now for AI. There's…AI vendors…claiming accuracy out in the field of a certain amount…It's incredible that you end up purchasing something…and trust without any real data on its performance…we're going to…[need] a relatively neutral repository of real data that people have collected on different algorithms." (Radiologist, Academic Medical Center, US)

Another US radiologist called for 'public summaries' of performance, while some UK participants lamented the lack of communication across different NHS trusts. In both the US and UK, liability concerns and professional reputation were raised as potential barriers to sharing this information.

Stakeholder Engagement

Participants highlighted the need for engaging stakeholders at every stage of AI deployment. A UK radiologist suggested that the lack of stakeholder engagement before introducing a stroke detection algorithm had made practitioners biassed against using it, with some even turning the system off (Radiologist, UK, AHSC). A few radiation oncologists (in both the US and UK) emphasized their scientific expertise—adding that if they were expected to give advice on quality assurance, it was only fair if they were involved with procurement as well. A UK radiation oncologist mentioned that accreditations might be helpful to ensure that there are designated AI experts at each NHS trust responsible for overseeing AI deployments (Clinical Scientist, UK, non-AHSC).

Other Participant Recommendations

Participants occasionally mentioned that clinical trials or prospective studies should be conducted for AI devices. One participant suggested that companies should subsidize trials, citing the example of clinical trials for drugs (Clinical Scientist, UK, non-AHSC). Less commonly, participants suggested that patients should be informed if AI is being used in their care, through obtaining explicit consent (Radiologist, UK, non-AHSC), or being informed by discharge summary (Medical Physicist, US, Academic Medical Center). A few UK-specific suggestions were also raised by radiation oncologists: one participant mentioned the need for overhauling information governance, while another participant highlighted the need for new AI-specific NHS procurement guidelines (Clinical Scientist, UK, non-AHSC).

Discussion

Our interview findings underscore clear challenges emerging in responsibility attribution and governance of AI in radiology. Participants highlighted the importance of resources and expertise in determining the nature of both AI

adoption and governance, echoing previous studies (Baten 2024; Raimo et al. 2023; Zhang et al. 2013; Chow, Lee, and Wu 2025). However, a newer insight was participants' reflections on how the perception of AI as 'just another piece of software' versus a formal medical device may be impacting the way that AI is adopted and governed as well.

Although participants generally viewed clinicians as having "final responsibility" over AI-related outcomes, echoing previous findings (Jones, Thornton, and Wyatt; Van Cauwenberge et al. 2022), participants noted that increasingly autonomous deployments of AI might challenge this existing framework, especially if errors were due to a lack of due diligence by commercial vendors. Some highlighted already shifting responsibility norms in instances where AI had been used to deal with worker shortages after hours. At the same time, participants commonly noted that staff in radiology departments were not adequately trained to use AI systems. These findings suggest that the 'moral crumple zone' presents as a real risk for healthcare practitioners.

When it comes to governance, it is unclear who should be responsible for the various different quality control measures necessitated by AI systems, including pre-deployment testing and post-deployment monitoring. This lack of clarity stems not only from the nascent stage of literature and best practices on AI evaluations, but also the lack of capacity for most health systems to take on such a task. The most rigorous AI governance committees tended to take on this role, but other committees relied on individual 'clinical champions' or those interested in purchasing the AI system, to monitor performance and report back to the committee. More centralized approaches to performance monitoring, such as an NHS-wide initiative, or third-party evaluations were also suggested. It is possible that such centralized approaches might even be more effective at catching population-wide impacts of these systems, but it is unclear if a centralized body will take on this responsibility. Here, a tension emerges between top-down and bottom-up approaches to governance.

Similarly, the profession also faces challenges in deciding who should be responsible for issuing guidelines and sharing best practices for AI. Participants highlighted concerns about too many professional and regulatory bodies being involved in AI ('many hands'), sometimes having competing priorities/goals. To that end, participants commonly mentioned a desire for better coordination and more streamlined guidelines. When conducting further research, we came across one cross-society statement from US professional societies on AI (Brady et al. 2024). Whether or not participants were aware of this statement, it appeared as though they were looking for stronger, implementable guidelines. However, participants also noted that the liability norms in medicine could be a barrier to knowledge sharing and strong 'guidelines', as individuals and institutions typically wait to establish a strong evidence base before putting out information. But establishing this evidence base for AI is tricky, as adoption is happening much quicker than people can study its effects, and in countries like the UK, information governance often prevents research groups from being able to access data for years at a time, slowing down the process of evaluation. Moreover, the performance of AI systems can be very dependent on characteristics of particular patient populations, making it that much more challenging to provide 'general' recommendations.

Motivated groups of individuals have stepped into this context to design novel AI-specific policies and approaches at their individual institutions. These 'governance entrepreneurs' commit resources to developing new governance interventions that they believe will pay off in the long run. However, like entrepreneurs in other contexts, they can only be successful if they have access to adequate resources, such as capital and expert input. Our findings suggest that lower resource hospitals tend to lack the resources to adequately test and monitor AI system performance. Therefore, some hospitals are relying on temporary or voluntary employees, such as medical residents, for testing and validation purposes. Yet this does not offer a sustainable path for ensuring the safety of AI systems over time. Thus, the current manner of AI deployment risks entrenching existing inequalities across health systems.

Importantly, the governance entrepreneurs we interviewed differ from entrepreneurs in other fields as they continue to operate within deeply hierarchical organizations. Thus, they must navigate power relations, justify their role, and rely on buy-in from higher ups on the importance of AI-specific governance. Many participants who had pushed for stronger AI governance were radiation oncologists. Some of them mentioned that despite their scientific expertise, and responsibility for quality assurance, they were sometimes overlooked when AI systems were deployed at their hospital. The power held by an AI governance committee is also highly organizationally specific, with some having strict enforcement power and others merely guiding or overseeing AI deployments. Additionally, the medical profession's intentionally siloed, specialty-driven approach appears to be at-odds with increasingly wide-ranging use of AI across different contexts. This creates new challenges in understanding whose expertise should be included when overseeing AI deployments, and how wide the scope of governance committees should be.

The governance solutions suggested by participants require resources, navigating political tensions, and breaking down legal and structural barriers to collaboration. Developing better AI training for staff would require resources to be invested either at the individual hospital-level, or on a higher level, by institutions involved in setting medical school curricula. Incentivizing individual hospitals to be

more transparent about AI use and performance would require resources for monitoring and publishing this information, and also potentially require addressing liability issues. Coordination across individual hospitals and professional bodies could require a restructuring of the specialty-driven approach to medical knowledge sharing, and buy-in from professional and regulatory bodies with complex political entanglements. Participants' suggestions of more streamlined 'general' guidance, while also calling for better local governance infrastructure, also present potential tensions between top-down and bottom-up approaches to governance. General guidelines might be less desired if they are made mandatory and prevent local hospitals from accounting for their own unique patient populations. On the other hand, leaving governance completely up to individual hospitals could create equity issues.

Policy interventions could be pursued by governing bodies at various levels. Governments could create opportunities for collaboration and streamlining of AI-related guidance. Grants could be made available for hospitals lacking the resources to set up their own governance committees. Hospitals could also be encouraged to consider whether they have the capacity for conducting due diligence prior to implementing AI solutions. Independent, umbrella organizations could be created to enable institutions to share information, and research funding could be made available for generating real-world evidence on AI deployments. Centralized approaches to performance monitoring could also be pursued, and potentially subsidized by commercial vendors.

The need for resources poses a real challenge to implementing the interventions suggested above. Investing in AI governance would make AI implementation more expensive, and probably be seen as counterintuitive to the 'efficiency' goals that hospital systems hope that AI will help them achieve. However, poorly implemented AI systems run the risk of either being shut off by clinicians, making the investment a waste, or threatening the well-being of patients and clinicians alike. While some might argue that the eventual efficiency gains realized by AI will pay for initial investments, this is not yet clear. As some industry and regulatory participants pointed out, it is hard to place a price on the savings that AI will bring in the future.

In future work, it would be interesting to explore the practical viability of these different interventions. Placing these recommendations within the broader context of current AI-related policy instruments, like healthcare AI regulatory sandboxes (Qiu et al. 2025), could also enable more reflection on the relationship between bottom-up and top-down approaches. With a greater sample size of participants across different countries, more explicitly comparative analysis could be done across jurisdictions. Researchers could also extend this work to other medical specialty areas, beyond radiology, and other domains, like government AI adoption, to see if similar trends emerge.

Limitations

Our study presents the following limitations. First, our lack of access to healthcare professionals from lower-resource hospitals limits our equity-related analysis to participant perspectives that were derived from anecdotal conversations with other colleagues. Second, our lack of inclusion of IT professionals makes it possible that we missed context regarding the role of hospital-level infrastructure challenges in AI adoption and governance. Additionally, being a single-coded study, our interpretation of the results may contain biases, although these were mitigated through regular peer debriefing. Finally, although our sample size of 22 participants reached saturation, given the qualitative nature of this study, the conclusions presented are not necessarily generalizable.

Conclusion

In this paper, we explore the emergence of 'bottom-up' approaches to AI governance in radiology departments in the US and UK. In interviewing 22 experts in clinical adoption and governance of AI in radiology, we survey the current state of AI governance in hospitals. The findings of this study highlight novel challenges to responsibility and governance generated by AI, including the assignation of responsibility for new governance functions like post-deployment monitoring, and re-negotiation of responsibility norms when AI is deployed more autonomously. Additionally, differing views on which governance functions should be conducted locally versus centrally indicate an underlying tension between bottom-up and top-down approaches. In our in-depth examination of committees of 'governance entrepreneurs' within different hospitals, we find great variation in the scope, composition, remit, and role of these committees, and note challenges faced in navigating power relations and resource constraints within hierarchical organizations. In particular, we note that the medical profession's specialty-driven approach appears increasingly at-odds with the wide-ranging use of AI across different contexts in the hospital, creating new challenges in understanding whose expertise should be included when overseeing AI deployments. While professional societies and governing bodies have released guidelines to help hospitals govern AI, the overwhelming number of competing recommendations has made them difficult to implement. We conclude with some recommended interventions to cope with these challenges, including investing resources into staff training and local governance, and improving coordination across siloed institutions.

Acknowledgments

The authors would like to first thank the participants of this study, without whom this study would not be possible. We would also like to thank Sam Bennett for their invaluable feedback on an earlier draft of this manuscript, Reshma Munbodh for her early guidance on this project, and Kathrin Cresswell for her feedback on research design. Additionally, we thank members of Purdue University's Governance and Responsible AI Lab (GRAIL) and Centre for Technomoral Futures for their feedback during work in progress seminars. Bhargavi Ganesh's PhD scholarship is sponsored by the Baillie Gifford gift to the Edinburgh Futures Institute to research the ethics of data and artificial intelligence at the Centre for Technomoral Futures.

Ethics Statement

The research protocol was reviewed by an ethics committee at the University of Edinburgh's School of Informatics. Participants were provided with a standard participant information sheet and consent form prior to each interview. Participants could consent to either audio or video recording, and were informed that their participation was voluntary, with the option to withdraw at any time. They were also made aware of how their interview data would be stored and retained for the duration of the study. All participant data was anonymized at the point of transcription, with individual participants assigned a random generated Participant ID number. Furthermore, participant affiliations were kept confidential to avoid identifiability. Due to the aforementioned steps, we do not foresee any risks for participants of this study.

Researcher Positionality

The authors of this study are based at academic institutions in the United Kingdom and United States. The main researcher has a background as a policy practitioner and researcher at the intersection of AI ethics, governance and responsibility, while the two co-authors have academic backgrounds researching AI policy and sociotechnical health systems, respectively. These backgrounds impacted the framing and analysis of the qualitative study.

References

Abdelwanis, M.; Alarafati, H. K.; Tammam, M. M. S; and Simsekler, M. C. E. 2024. Exploring the Risks of Automation Bias in Healthcare Artificial Intelligence Applications: A Bowtie Analysis. *Journal of Safety Science and Resilience* 5 (4): 460–69. https://doi.org/10.1016/j.jnlssr.2024.06.001.

Abimbola, S.; Negin J.; Martiniuk, A. L.; and Jan, S. 2017. Institutional Analysis of Health System Governance. *Health Policy and Planning* 32 (9): 1337–44. https://doi.org/10.1093/heapol/czx083.

Adam, H.; Yang, M. Y.; Cato, K.; Balidini, I.; Senteio, C.; Celi, L. A.; Zeng, J.; Singh, M.; and Ghassemi, M. 2022. Write It Like You See It: Detectable Differences in Clinical Notes by Race Lead to Differential Model Recommendations. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society,* AIES '22, 7-21. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/3514094.3534203.

Allen, B.; Agarwal S.; Coombs L.; Wald, C.; and Dreyer, K. 2021. 2020 ACR Data Science Institute Artificial Intelligence Survey. *Journal of the American College of Radiology* 18 (8): 1153–59. https://doi.org/10.1016/j.jacr.2021.04.002.

American Medical Association. 2024. Augmented Intelligence Development, Deployment, and Use in Health Care. https://www.ama-assn.org/system/files/ama-ai-principles.pdf. Accessed: 2025-04-04.

Baten, R. B. A. 2024. How Are US Hospitals Adopting Artificial Intelligence? Early Evidence from 2022. *Health Affairs Scholar* 2 (10): qxae123. https://doi.org/10.1093/haschl/qxae123.

Beauchamp, T. L., and Childress, J. F. 1979. Principles of biomedical ethics. Oxford: Oxford University Press.

Bleher, H., and Braun, M. 2022. Diffused responsibility: attributions of responsibility in the use of AI-driven clinical decision support systems. *AI and Ethics* 2 (4):747-761. https://doi.org/10.1007/s43681-022-00135-x.

Brady, A. P.; Allen, B.; Chong, J.; Kotter, E.; Kottler, N.; Mongan, J.; Oakden-Rayner, L.; Dos Santos, D. P.; Tang, A.; Wald, C.; and Slavotinek, J. 2024. Developing, Purchasing, Implementing and Monitoring AI Tools in Radiology: Practical Considerations. A Multi-Society Statement From the ACR, CAR, ESR, RANZCR & RSNA. *Journal of the American College of Radiology* 21 (8): 1292–310. https://doi.org/10.1016/j.jacr.2023.12.005.

British Medical Association. 2024. Principles for artificial intelligence (AI) and its application in healthcare. https://www.bma.org.uk/advice-and-support/nhs-delivery-and-workforce/technology/principles-for-artificial-intelligence-ai-and-its-application-in-healthcare. Accessed: 2025-04-04.

Bryden, D., and Storey, I. 2011. Duty of care and medical negligence. *Continuing Education in Anaesthesia Critical Care & Pain* 11 (4): 124-127. https://doi.org/10.1093/bjaceaccp/mkr016.

Bogner, A.; Littig, B.; and Menz, W. 2009. *Interviewing Experts*. London: Palgrave Macmillan UK.

Castonguay, A.; Wagner, G.; Motulsky, A.; and Paré, G. 2024. 'AI Maturity in Health Care: An Overview of 10 OECD Countries'. *Health Policy* 140 (February): 104938. https://doi.org/10.1016/j.healthpol.2023.104938.

Chen, S.; Guevara, M.; Moningi, S.; Hoebers, F.; Elhalawani, H.; Kann, B. H.; Chipidza, F. E.; Leeman, J.; Aerts, H. J. W. L.; Miller, T.; Savova, G. K.; Gallifant, J.; Celi, L. A. Mak, R. H.; Lustberg, M.; Afshar, M.; and Bitterman, D. S. 2024. The Effect of Using a Large Language Model to Respond to Patient Messages. *The Lancet Digital Health* 6 (6): e379–81. https://doi.org/10.1016/S2589-7500(24)00060-8.

Cobbe, J.; Veale, M.; and Singh, J. Understanding accountability in algorithmic supply chains. 2023. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '23: 1186–1197. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/3593013.3594073.

Coeckelbergh, M. 2020. Artificial Intelligence, Responsibility Attribution, and a Relational Justification of Explainability. *Science*

and Engineering Ethics 26 (4): 2051-2068. https://doi.org/10.1007/s11948-019-00146-8.

Chow, J.; Lee, R.; and Wu, H. 2025. How Do Radiologists Currently Monitor AI in Radiology and What Challenges Do They Face? An Interview Study and Qualitative Analysis. *Journal of Imaging Informatics in Medicine*, ahead of print, April 8. https://doi.org/10.1007/s10278-025-01493-8.

Daye, D.; Wiggins, W. F.; Lungren, M. P.; Alkasab, T.; Kottler, N.; Allen, B.; Roth, C. J.; Bizzo, B. C.; Durniak, K.; Brink, J. A.; Larson, D. B.; Dreyer, K. J.; and Langlotz, C. P. 2022. Implementation of Clinical Artificial Intelligence in Radiology: Who Decides and How? *Radiology* 305 (3): 555–63. https://doi.org/10.1148/radiol.212151.

Denzin, N. K. *The Research Act: A Theoretical Introduction to Sociological Methods.* 2017. Routledge. https://doi.org/10.4324/9781315134543.

Dratsch, T.; Chen, X.; Mehrizi, M.; Kloeckner, R.; Mähringer-Kunz, A.; Püsken, M.; Baeßler, B.; Sauer, S.; Maintz, D.; and Pinto dos Santos, D. 2023. Automation Bias in Mammography: The Impact of Artificial Intelligence BI-RADS Suggestions on Reader Performance. *Radiology* 307 (4): e222176. https://doi.org/10.1148/radiol.222176.

Durant, R. F., and Diehl, P. F. 1989. Agendas, Alternatives, and Public Policy: Lessons from the U.S. Foreign Policy Arena. *Journal of Public Policy* 9 (2): 179–205.

Elish, M. C. 2019. Moral Crumple Zones: Cautionary Tales in Human-Robot Interaction. *Engaging Science, Technology, and Society* 5 (2019): 40–60. https://doi.org/10.17351/ests2019.260.

Faruki, A. A.; Zane, R. D.; and Wiler, J. L. 2022. The Role of Academic Health Systems in Leading the "Third Wave" of Digital Health Innovation. *JMIR Medical Education* 8 (4): e32679. https://doi.org/10.2196/32679.

Fazakarley, C. A.; Breen, M.; Leeson, P.; Thompson, B.; and Williamson, V. 2023. Experiences of Using Artificial Intelligence in Healthcare: A Qualitative Study of UK Clinician and Key Stakeholder Perspectives. *BMJ Open* 13 (12): e076950. https://doi.org/10.1136/bmjopen-2023-076950.

Ford, E.; Edelman, N.; Somers, L; Shrewsbury, D.; Lopez Levy M.; van Marwijk H.; Curcin V.; and Porat T. 2021. Barriers and Facilitators to the Adoption of Electronic Clinical Decision Support Systems: A Qualitative Interview Study with UK General Practitioners. *BMC Medical Informatics and Decision Making* 21 (1): 193. https://doi.org/10.1186/s12911-021-01557-z.

Foy, R.; Skrypak, M,; Alderson, S.; Ivers, N. M.; McInerney, B.; Stoddart, J.; Ingham, J.; and Keenan, D. Revitalising audit and feedback to improve patient care. *BMJ*. 368: m213. https://doi.org/10.1136/bmj.m213.

Garcia, P.; Ma, S. P.; Shah, S.; Smith, M.; Jeong, Y.; Devon-Sand, A.; Tai-Seale, M.; Takazawa, K.; Clutter, D.; Vogt, K.; Lugtu C.; Rojo M.; Lin, S. Shanafelt, T.; Pfeffer, M. A.; and Sharp, C. 2024. Artificial Intelligence–Generated Draft Replies to Patient Inbox Messages. *JAMA Network Open* 7 (3): e243201. https://doi.org/10.1001/jamanetworkopen.2024.3201.

Garg, A. X.; Adhikari, N. K. J; McDonald, H.; Rosas-Arellano, M. P.; Devereaux, P. J; Beyene, J.; Sam, J.; and Haynes, B. R. 2005. Effects of Computerized Clinical Decision Support Systems on Practitioner Performance and Patient Outcomes: A Systematic Review. *JAMA* 293 (10): 1223–38. https://doi.org/10.1001/jama.293.10.1223.

Gichoya, J. W.; Banerjee, I.; Bhimireddy, A. R.; Burns, J. L.; Celi, L. A.; Chen, L.; Correa, R.; Dullerud, N.; Ghassemi, M.; Huang, S-C; Kuo, P-C; Lungren, M. P.; Palmer, L. J.; Price, B. J.; Purkayastha, S.; Pyrros, A. T.; Oakden-Rayner, L.; Okechukwu, C; Seyyed-Kalantari, L.; Trivedi, H.; Wang, R.; Zaiman, Z.; and Zhang, H. AI Recognition of Patient Race in Medical Imaging: A Modelling Study. *The Lancet Digital Health* 4 (6): e406–14. https://doi.org/10.1016/S2589-7500(22)00063-2.

Gilbert, S.; Harvey, H.; Melvin, T.; Vollebregt, E.; and Wicks, P. 2023. Large Language Model AI Chatbots Require Approval as Medical Devices. *Nature Medicine* 29 (10): 2396–98. https://doi.org/10.1038/s41591-023-02412-6.

GOV.UK. 2023. £21 Million to Roll out Artificial Intelligence across the NHS. https://www.gov.uk/government/news/21-million-to-roll-out-artificial-intelligence-across-the-nhs. Accessed: 2025-08-30.

Guba, E. G. 1981. Criteria for assessing the trustworthiness of naturalistic inquiries. *ECTJ* 29: 75-91. https://doi.org/10.1007/BF02766777.

Hadjiiski, L.; Cha, K.; Chan, H.-P.; Drukker, K.; Morra, L.; Näppi, J. J.; Sahiner, B.; Yoshida, H.; Chen, Q.; Deserno, T. M.; Greenspan, H.; Huisman, H.; Huo, Z.; Mazurchuk, R.; Petrick, N.; Regge, D.; Samala, R.; Summers, R. M.; Suzuki, K.; Tourassi, G.; Vergara, D.; and Armato III, S. G. 2023. AAPM task group report 273: Recommendations on best practices for AI and machine learning for computer-aided diagnosis in medical imaging. *Medical Physics* 50 (2): e1-e24. https://doi.org/10.1002/mp.16188.

Hajar, R. 2017. The Physician's Oath: Historical Perspectives. *Heart Views: The Official Journal of the Gulf Heart Association* 18 (4): 154-159. https://doi.org/10.4103/HEARTVIEWS.HEARTVIEWS_131_17.

Esposito, P., and Dal Canton, A. 2014. Clinical audit, a valuable tool to improve quality of care: General methodology and applications in nephrology. *World Journal of Nephrology* 3 (4): 259-255. https://doi.org/10.5527/wjn.v3.i4.249.

Jalilvand, M. A.; Raeisi, A. R.; and Shaarbafchizadeh, N. 2024. Hospital governance accountability structure: a scoping review. *BMC Health Services Research* 24: 47. https://doi.org/10.1186/s12913-023-10135-0.

Jones, C.; Thornton, J.; and Wyatt, J. C. 2023. Artificial Intelligence and Clinical Decision Support: Clinicians' Perspectives on Trust, Trustworthiness, and Liability. *Medical Law Review* 31 (4): 501–20. https://doi.org/10.1093/medlaw/fwad013.

Kim, J. Y.; Boag, W.; Gulamali, F.; Hasan, A.; Hogg, H. D. J; Lifson, M.; Mulligan, D.; Patel, M.; Raji, I. D.; Sehgal, A.; Shaw, K.; Tobey, D.; Valladares, A.; Vidal, D.; Balu, S.; and Sendak, M. 2023. Organizational Governance of Emerging Technologies: AI Adoption in Healthcare. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency,* FAccT '23: 1396–1417. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/3593013.3594089.

Knight, D. R. T.; Aakre, C. A.; Anstine, C. V.; Munipalli, B.; Biazar, P.; Mitri, G.; Valery, J. R.; Brigham, T.; Niazi, S. K.; Perlman, A. I.; Halamka, J. D.; and Dabrh, A. M. A. Artificial Intelligence for Patient Scheduling in the Real-World Health Care Setting: A Metanarrative Review. *Health Policy and Technology* 12 (4): 100824. https://doi.org/10.1016/j.hlpt.2023.100824.

Krishnamoorthy, M.; Sjoding, M. W.; and Wiens, J. 2024. Off-label use of artificial intelligence models in healthcare. 2024. *Nature Medicine* 30 (6): 1525-1527. https://doi.org/10.1038/s41591-024-02870-6.

Kwee, T. C., and Kwee, R. M. 2021. Workload of Diagnostic Radiologists in the Foreseeable Future Based on Recent Scientific Advances: Growth Expectations and Role of Artificial Intelligence. *Insights into Imaging* 12 (1): 88. https://doi.org/10.1186/s13244-021-01031-4.

Laï, M.-C.; Brian, M.; and Mamzer, M.-F. 2020. Perceptions of Artificial Intelligence in Healthcare: Findings from a Qualitative Survey Study among Actors in France. *Journal of Translational Medicine* 18 (1): 14. https://doi.org/10.1186/s12967-019-02204-y.

Liberati, E. G.; Ruggiero, F.; Galuppo, L.; Gorli, M.; González-Lorenzo, M.; Maraldi, M.; Ruggieri, P.; Polo Friz, H.; Scaratti, G.; Kwag, K. H.; Vespignani, R.; and Moja, L. 2017. What Hinders the Uptake of Computerized Decision Support Systems in Hospitals? A Qualitative Study and Framework for Implementation. Implementation *Science* 12 (1): 113. https://doi.org/10.1186/s13012-017-0644-2.

Liberatore, M. J., and Nydick, R. L. 2008. The analytic hierarchy process in medical and health care decision making: A literature review. *European Journal of Operational Research* 189 (1): 194-207. https://doi.org/10.1016/j.ejor.2007.05.001.

Lincoln, Y, S.; Guba, E. G.; and Pilotta, J. J. 1985. Naturalistic Inquiry. *International Journal of Intercultural Relations* 9 (4): 438–39. https://doi.org/10.1016/0147-1767(85)90062-8.

Matthias, A. 2004. The Responsibility Gap: Ascribing Responsibility for the Actions of Learning Automata. *Ethics and Information Technology* 6 (3): 175–83. https://doi.org/10.1007/s10676-004-3422-1.

Morse, J. M.; Knafl, K. A.; and Breitmayer, B. J. 1991. Triangulation in Qualitative Research: Issues of Conceptual Clarity and Purpose. *Qualitative Nursing Research: A Contemporary Dialogue*. SAGE Publications, Inc. https://doi.org/10.4135/9781483349015.

Murdoch, B. 2021. Privacy and Artificial Intelligence: Challenges for Protecting Health Information in a New Era. *BMC Medical Ethics* 22 (1): 122. https://doi.org/10.1186/s12910-021-00687-3.

Nandi, P. L. 2000. Ethical Aspects of Clinical Practice. *Archives of Surgery* 135 (1): 22-25. https://doi.org/10.1001/archsurg.135.1.22

Nissenbaum, H. 1996. Accountability in a computerized society. *Science and Engineering Ethics* 2 (1): 25-42. https://doi.org/10.1007/BF02639315.

Nittas, V.; Ormond, K. E.; Vayena, E.; and Blasimme, A. 2025. Realizing the Promise of Machine Learning in Precision Oncology: Expert Perspectives on Opportunities and Challenges. *BMC Cancer* 25 (1): 276. https://doi.org/10.1186/s12885-025-13621-2.

Nong, P.; Hamasha, R.; and Platt, J. 2024. Equity and AI Governance at Academic Medical Centers. *The American Journal of Managed Care* 30 (Spec Issue No. 6): SP468–72. https://doi.org/10.37765/ajmc.2024.89555.

Nyholm, S. 2018. Attributing Agency to Automated Systems: Reflections on Human–Robot Collaborations and Responsibility-Loci. *Science and Engineering Ethics* 24 (4): 1201-1219. https://doi.org/10.1007/s11948-017-9943-x.

Ong, J. C. L.; Ning, Y.; Liu, M.; Ma, Y.; Liang, Z.; Singh, K.; Chang, R.; Vogel, S.; Lim, J. C. W.; Tan, I. S. K.; Freyer, O.; Gilbert, S.; Bitterman, D. S.; Liu, X.; Denniston, A. K.; and Liu, N. 2025. Regulatory Science Innovation for Generative AI and Large Language Models in Health and Medicine: A Global Call for Action. arXiv:2502.07794. Preprint, arXiv, January 27. https://doi.org/10.48550/arXiv.2502.07794.

Oremule, B.; Saunders, G. H.; Kluk, K.; d'Elia, A.; and Bruce, I. A. 2024. Understanding, Experience and Attitudes towards Artificial Intelligence Technologies for Clinical Decision Support in Hearing Health: A Mixed-Methods Survey of Healthcare Professionals in the UK. *The Journal of Laryngology & Otology* 138 (9): 928–35. https://doi.org/10.1017/S0022215124000550.

Palaniappan, K.; Lin, E. Y. T.; and Vogel, S. 2024. Global Regulatory Frameworks for the Use of Artificial Intelligence (AI) in the Healthcare Services Sector. *Healthcare* 12 (5): 562. https://doi.org/10.3390/healthcare12050562.

Parikh, R. B.; Manz, C. R.; Nelson, M. N.; Evans, C. N.; Regli, S. H.; O'Connor, N.; Schuchter, L. M.; Schulman, L. N.; Patel, M. S.; Paladino, J.; and Shea, J. A. 2022. Clinician Perspectives on Machine Learning Prognostic Algorithms in the Routine Care of Patients with Cancer: A Qualitative Study. *Supportive Care in Cancer* 30 (5): 4363–72. https://doi.org/10.1007/s00520-021-06774-w.

Pope, C.; Halford, S.; Turnbull, J.; Prichard, J.; Calestani, M.; and May, C. 2013. Using Computer Decision Support Systems in NHS Emergency and Urgent Care: Ethnographic Study Using Normalisation Process Theory. *BMC Health Services Research* 13 (1): 111. https://doi.org/10.1186/1472-6963-13-111.

Raimo, N.; De Turi, I.; Albergo, F.; and Vitolla, F. 2023. The Drivers of the Digital Transformation in the Healthcare Industry: An Empirical Analysis in Italian Hospitals. *Technovation* 121 (March 2023): 102558. https://doi.org/10.1016/j.technovation.2022.102558.

Sahiner, B.; Chen, W.; Samala, R. K.; and Petrick, N. 2023. Data drift in medical machine learning: implications and potential remedies. *The British Journal of Radiology* 96 (1150). https://doi.org/10.1259/bjr.20220878.

Santoni de Sio, F., and Mecacci, G. 2021. Four Responsibility Gaps with Artificial Intelligence: Why they Matter and How to Address them. *Philosophy and Technology* 34 (4): 1057-1084. https://doi.org/10.1007/s13347-021-00450-x.

Thompson, J. 2022. A Guide to Abductive Thematic Analysis. *The Qualitative Report*, 27(5): 1410-1421. https://doi.org/10.46743/2160-3715/2022.5340

Smith, H., and Fotheringham, K. 2020. Artificial Intelligence in Clinical Decision-Making: Rethinking Liability. *Medical Law International* 20 (2): 131–54. https://doi.org/10.1177/0968533220945766.

The Royal College of Radiologists. 2024. AI deployment fundamentals for medical imaging. https://www.rcr.ac.uk/our-services/all-our-publications/clinical-radiology-publications/ai-deployment-fundamentals-for-medical-imaging/. Accessed: 2025-04-04.

US Food and Drug Administration. 2025. Artificial Intelligence and Machine Learning (AI/ML)-Enabled Medical Devices. https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices. Accessed: 2025-08-31.

Van Cauwenberge, D.; Van Biesen, W.; Decruyenaere, J.; Leune, T.; and Sterckx, S. 2022. "Many Roads Lead to Rome and the Artificial Intelligence Only Shows Me One Road": An Interview Study on Physician Attitudes Regarding the Implementation of Computerised Clinical Decision Support Systems. *BMC Medical Ethics* 23 (1): 50. https://doi.org/10.1186/s12910-022-00787-8.

Vallor, S., and Ganesh, B. 2023. Artificial intelligence and the imperative of responsibility. In *The Routledge Handbook of Philosophy of Responsibility*, edited by M. Kiener, 395-406. London: Routledge.

Van Veen, D.; Van Uden, C.; Blankemeier, L.; Delbrouck J.-B.; Aali, A.; Bluethgen, C.; Pareek, A.; Polacin, M.; Reis, E. P.; Seehofnerová, A.; Rohatgi, N.; Hosamani, P; Collins, W.; Ahuja, N.; Langlotz, C. P.; Hom, J.; Gatidis, S.; Pauly, J.; and Chaudhari, A. S. 2024. Adapted Large Language Models Can Outperform Medical Experts in Clinical Text Summarization. *Nature Medicine* 30 (4): 1134–42. https://doi.org/10.1038/s41591-024-02855-5.

Vogt, P. W., and Johnson, R. B. 2015. *The SAGE Dictionary of Statistics & Methodology: A Nontechnical Guide for the Social Sciences*. SAGE Publications.

Widder, D. G., and Nafus, D. 2023. Dislocated accountabilities in the "AI supply chain": Modularity and developers' notions of responsibility. *Big Data & Society* 10 (1). https://doi.org/10.1177/20539517231177620.

Qiu, Y.; Yao, H.; Tian, X.; and You, M. 2025. Regulatory sandbox expansion: Exploring the leap from fintech to medical artificial intelligence. *Intelligent Oncology* 1 (2): 120-127. https://doi.org/10.1016/j.intonc.2025.03.001.

Zhang, N. J.; Seblega, B.; Wan, T.; Unruh, L.; Agiro, A.; and Miao, L. 2013. Health Information Technology Adoption in U.S. Acute Care Hospitals. *Journal of Medical Systems* 37 (2): 9907. https://doi.org/10.1007/s10916-012-9907-2.