Frugal and Agentic AI for Medical Imaging in Resource-Constrained Settings: A Sustainable Framework

Roseline Oluwaseun Ogundokun[1,2, 3], Rotimi-Williams Bello[1], Pius Adewale Owolawi[1]

[1] Department of Computer Systems Engineering, Faculty of Information and Communication Technology, Tshwane University of Technology, Pretoria, South Africa.
[2] Department of Multimedia Engineering, Kaunas University of Technology, Kaunas, Lithuania
[3] Department of Computer Science, Redeemer's University, Ede, Osun State, Nigeria.

Ogundokunroseline1@gmail.com, roseogundokun@gmail.com

Abstract. Healthcare systems with limited resources face critical shortages of imaging capabilities, which compound diagnostic disparities. While artificial intelligence is a potential solution, the data and technology required to implement typical AI models make them impractical in many resource-limited medical facilities. This paper describes the theoretical framework for the resource-aware, agentically frugal intelligent (RAFI) development of medical imaging Artificial Intelligence (AI), with the goal of developing efficient, independent, and just medical imaging AI systems in low-resource settings. RAFI combines frugal machine learning techniques, agentic decision-making, and trustworthy concepts to help overcome data scarcity, computational limitations, and bias. We synthesise the literature and identify the main gaps. The RAFI framework contains a list of principles (e.g., resource prioritization, fairness by design, and continuous local adaptation) and sets objectives based on resource and ethical constraints. A comparison between RAFI and current paradigms highlights differences in scalability, equity, and sustainability. We provide four hypotheses (P1-P4) for future research; we present a roadmap for the implementation and validation of RAFI; and we discuss the implications for policy and capacity building. The development of RAFI provides a means to build a scalable, trustworthy AI system for equitable global health.

Keywords: Healthcare system, Agentic AI, Medical imaging, Resource-constrained settings, Sustainability

1 Introduction

Deep learning has led to significant improvements in how we analyze medical images. However, success depends on the availability of large amounts of data, computing resources, and infrastructure [1, 2]. In many cases, current best-in-class models have millions of parameters and need GPU clusters for both training and inference [3, 4]. Computational scale has limited our ability to make progress solely through accuracy improvements, while also leading to extremely high deployment costs and increased carbon emissions. Recent work has shown that some of the largest AI models can emit more carbon than multiple vehicles over their lifetimes [3]. Within a healthcare context, especially when every site of care has significant constraints to operate from (e.g., no access to even a single MRI or GPU), these issues of "Green AI" are much more pronounced, as many low-resource or rural hospitals operate without reliable access to either power or internet connectivity, making it nearly impossible for them to do inference in the cloud [5, 6].

Furthermore, the common AI deployment contradiction is that even though they are very efficiently designed, most models do not provide benefits in real-life settings [7]; high privacy regulations along with the nature of medical images (e.g. easy re-identification via MRI) prohibit naive cloud processing [8], low-latency requirements (e.g. real time surgical assistance) restrict remote computing, and inequities within the infrastructure limits benefits of AI on a very localised basis to the well-funded centres. As a result of the foregoing, AI for medical images has often been limited to retrospective studies and pilot-type tools. Meanwhile, newer paradigms such as agentic AIs promise an active, workflow-optimising system; however, they have predominantly been designed and built for resource-rich environments [9, 10].

At present, algorithmic fairness and bias have been identified as critical concerns in AI technologies in medicine, including health care and medical imaging [11, 12]. Specifically, AI technologies trained on imbalanced (skewed) datasets tend to underperform on underrepresented groups (e.g., Minority groups) [11, 13, 14]. For instance, one study used a diabetic retinopathy model that achieved 73% accuracy in light-skinned patients and only 60.5% in dark-skinned patients [15]. When high- and low-income global organisations collaborate, unchecked algorithmic bias can exacerbate disparities in health equity within these populations [11, 12]. Furthermore, deploying AI in low-resource settings requires additional elements (i.e., infrastructure, training, governance processes, and policy support) beyond the technology itself [16].

The literature has advanced each of these areas independently (i.e., efficient models; agent autonomy; fairness; and global deployment of health-related AI) [4, 9-11, 16], but there is no single framework that integrates these areas for low-resource medical imaging. Reviews on model compression or Edge-AI typically do not address issues of workflow autonomy and fairness together, or vice versa [9, 10, 14]. This definitional gap obstructs the design of truly sustainable AI systems.

The Resource-Aware Agentic Frugal Intelligence (RAFI) framework will help integrate these principles. RAFI is a framework for designing, optimizing, and enhancing the use of medical imaging technologies; it has been designed to optimize three aspects of the system: Resource-Aware Frugality, Agentic Autonomy, and Trustworthiness & Equity. The main contributions of this paper are summarized as follows:

a. A novel conceptual framework (RAFI) of medical imaging AI specifically designed for resource-constrained environments.

b. A focused critical synthesis of literature on resource-efficient learning, agentic AI, fairness, trustworthy medical AI, and deployment in resource-constrained settings, highlighting the gaps that motivate RAFI.
c. Explicitly defined design principles for RAFI, including resource efficiency, fairness by design, agentic behaviour, and continual local adaptation.
d. A resource-constrained conceptual formulation and illustrative pseudocode showing how resource assessment, model efficiency, fairness evaluation, uncertainty handling, and bounded agentic actions can be incorporated within the RAFI lifecycle.
e. A conceptual comparison of RAFI to conventional AI approaches (see Table 2) showing anticipated advantages/disadvantages.
f. Testable propositions and practical evaluation considerations for future implementation and validation of RAFI.
g. A discussion of the policy, governance, workforce, and capacity-building considerations associated with applying RAFI in resource-constrained healthcare settings.
h. A set of open-ended research inquiries and hypotheses (P1–P4) intended to assist with directing the development of sustainable AI within the realm of medicine.

Our goal is to produce a rigorous, persuasive foundation that can inform both researchers and practitioners. By focusing on *frugality, agency, and trust* in unison, RAFI aims to enable AI-driven diagnostic tools that are not only accurate but also accessible, fair, and sustainable across diverse clinical environments.

RAFI's contribution is more than the mere creation of a new compression, fairness, or agentic algorithm. Rather, its contribution is organizing various components in a cohesive manner within a single resource-restricted medical imaging life cycle. In the case of RAFI, all aspects related to resource restrictions influence model choice and deployment; fairness is evaluated in the pre- and post-deployment phases; agentic actions are constrained by uncertainty and clinical supervision; and local adaptability remains under control and is subject to updates. This intent enables the framework to capture the interrelationships among resource efficiency, workflow support, equity, safety, and sustainability.

2 Related Work

This part of the document gives an organized narrative description of the materials connected to RAFI rather than a systematic review of them, and it centers on six areas of particular interest in medical imaging under resource constraints; these are: model efficiency, autonomy, equality and trust, data shift and domain shift, deployment infrastructures, and local capacity-building. There's been an interesting shift in how researchers have approached the construction of Medical AI Systems, moving away from focusing on accuracy as the primary metric of success toward developing frugal, efficient, and deployable systems. Essentially, Frugal Machine Learning focuses on developing models that deliver the same diagnostic performance while reducing the requirements for data, memory, computation, energy, and hardware. Several techniques for reducing the computational demands of models, such as pruning, quantization, knowledge distillation, lightweight architectures, and edge hardware, have been applied to Medical Imaging modalities, such as Portable Ultrasound and On-device Diagnostics. However, most lightweight models developed thus far have been evaluated only as stand-alone technical solutions and have not been evaluated against the clinical constraints that must be adhered to in clinical or diagnostic environments (such as Privacy, Latency, Infrastructure Heterogeneity, and Connectivity & Maintenance) [17-19].

Agentic AI has garnered interest in healthcare and radiology due to its potential to act proactively rather than merely react to user input. Some examples of possible applications of agentic systems in image analysis include facilitating the prioritization of urgent scans, reorganizing imaging workflows using machine learning, and providing context-adapted recommendations while also reducing healthcare provider workload. However, agentic systems are still in their infancy, with many unresolved issues regarding safety, explainability, regulation, and accountability, and existing agentic frameworks rarely address how autonomous systems may function in low-resource settings [9, 20, 21].

AI is a trustworthy and fair element of the 3rd pillar of Trust in AI. As encapsulated in the Future-AI framework, AI's characteristic of fairness can be assessed through the lens of universality, traceability, usability, robustness and explainability. WHO and UNESCO's policies, on the other hand, identify fairness as a human rights issue, the basis for consent, accountability, and equity. There are technically feasible ways to mitigate the impact of bias; however, in practice, fairness is typically an after-the-fact consideration rather than part of the design process. In addition, fairness-related effects of model compression have not been extensively investigated [22-24].

The use of AI in LMICs or other resource-limited environments is hindered by equipment shortages, unreliable power supply, poor connectivity and insufficiently trained personnel; evidence from RAD-AID continues to define the importance of infrastructure, education, phased implementation and policy support in relationship to AI; as such current research identifies efficiency, autonomy, fairness and contextual deployment as separate entities, which leads to the development of the RAFI Framework as a comprehensive framework to design medical imaging AI that is resource-aware, agentic, fair and sustainable [6, 25].

Table 1. Gap analysis of current approaches in medical imaging AI (columns indicate the state-of-the-art vs remaining research gaps).

Aspect	Current Approaches	Key Limitations / Gaps
Model Efficiency	Conventional deep learning models, including CNNs and Transformers, remain dominant; efficiency is commonly pursued through pruning, quantization, knowledge distillation, and lightweight architectures.	Many models remain computationally heavy; energy use, inference latency, memory demand, and deployment costs are rarely reported systematically. Most optimizations are evaluated in isolated experimental settings rather than within real clinical constraints.
Autonomy	Current systems mostly provide passive assistance for classification, segmentation, detection,	Agentic AI remains immature in medical imaging. Existing prototypes are rarely resource-aware, safety and human oversight mechanisms

	and triage; early agentic AI proposals suggest autonomous workflow management and clinical task planning.	remain underdeveloped, and integration with real imaging workflows is still insufficiently validated.
Fairness and Trust	Bias measurement, fairness-aware learning, explainable AI, and ethical guidelines such as WHO AI ethics guidance and FUTURE-AI are increasingly used to guide the development of trustworthy healthcare AI.	Fairness is often treated as a post hoc correction rather than a design principle. LMIC-specific population biases remain underexplored, and trustworthiness guidelines are rarely operationalized in real resource-constrained deployments.
Data and Domain Shift	Transfer learning, domain adaptation, federated learning, and privacy-preserving learning are used to address limited data, institutional variation, and privacy concerns.	LMIC datasets remain scarce and underrepresented. Models trained on high-income-country datasets may generalize poorly to local populations, scanners, disease patterns, and imaging protocols. Continual learning frameworks for field deployment are still limited.
Deployment	Cloud-based AI services, edge computing, mobile health applications, and early edge-AI imaging tools are being explored for clinical deployment.	Cloud-based systems may be restricted by privacy regulations, poor connectivity, latency requirements, and infrastructure limitations. Turnkey edge-AI systems for medical imaging remain limited, while maintenance, local technical support, and user training are often overlooked.
Capacity Building	Training programmes, radiology outreach initiatives, infrastructure-building projects, and collaborative implementation models support AI readiness in underserved settings.	Capacity-building efforts are often separated from technical AI design. Few frameworks integrate local workforce development, infrastructure planning, governance, and AI system development into a single deployment pathway.

RAFI is an iterative framework for resource-aware, agentic, and trustworthy medical imaging AI in resource-constrained settings. As shown in Figure 1, it begins with medical image and clinical context input, followed by quality and bias assessment to detect data weaknesses, imbalance, and unfairness. It then develops frugal models through compression, pruning, and efficient learning before introducing agentic decision-making for task selection, routing, and workflow support. Trust and safety checks assess fairness, explainability, privacy, and clinical oversight. The system is deployed locally at the edge, refined through clinician feedback, and continuously monitored for drift, performance, equity, and continual learning needs.

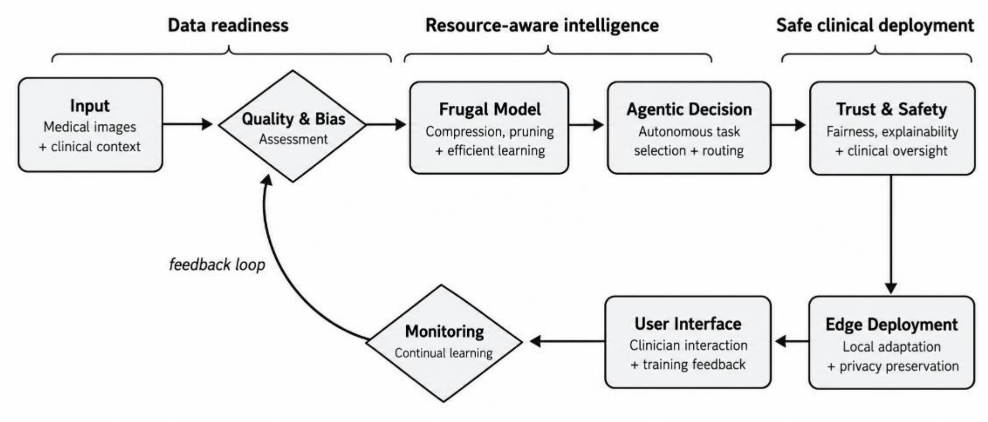

Fig. 1. RAFI workflow for resource-aware, agentic, and trustworthy medical imaging AI

Algorithmic Workflow (Pseudocode)

Algorithm I (i.e., the training and inference loop) sketch of how to implement RAFI in practice. This pseudo-code is visual and does not provide implementation details; however, it clearly illustrates how model frugality and fairness checks can be integrated into the training process and shows that the agentic layer can modify the output prior to a final decision.

1. Arduino
2. Copy
3. Algorithm 1: RAFI Training & Inference Workflow
4. Input: Training images X and labels Y, resource constraint profile $C = \{C_{latency}\ ,\ C_{memory}\ , C_{energy}\ , C_{storage}\ , C_{connectivity}\ , C_{cost}\}$, fairness threshold δ.
5. Initialize model θ_0 (lightweight architecture), agentic policy π_0.
6. For each epoch do:
7. for each minibatch (x,y) in (X,Y) do:
8. // Frugal adaptation
9. if Cost(θ) > C_max:
10. θ := compress(θ) *// prune/quantize to meet budget*
11. end if
12. *// Standard training step*
13. predictions = f_θ(x)
14. loss = CrossEntropy(predictions, y)
15. θ := UpdateParameters(θ, ∂loss)
16. end for
17. *// Fairness evaluation*
18. if FairnessGap(θ; X) > δ:
19. θ := applyDebiasing(θ, X) *// e.g. reweight, adversarial fine-tuning*
20. end if

```
end for
// Inference on new image x_new
features = preprocess(x_new)
output = f_θ(features)
if Uncertainty(output) > τ:
    // Agentic action: request additional data or analysis
    output = π(output, context)  // agent policy may select action or additional pipeline
end if
final_decision = output  // possibly after human review
return final_decision
```

In this formulation, A refers to information about the subgroup that can be considered ethically, legally, and clinically in the context of the planned use. The choice of fair measure would depend on the task and could mean that the differences in specificity, false-negative rate, calibration, sensitivity or any other subgroup indicator that can help in clinical work would be analyzed. One should not assume that the tolerance δ is universal; instead, it must be determined before starting the planned clinical task. When it is possible to identify an important difference, one might use resampling or reweighting the sample or apply any other technique suitable for this case.

Resource representation. In the context of RAFI, resource availability is understood as a set of site-specific constraints rather than a single, universally applicable limit. The resource profile can consist of maximum acceptable inference delay, available memory size, model capacity, energy availability, hardware configuration, connectivity, and throughput. The boundaries of these constraints are defined for a particular clinic prior to choosing or optimizing a model. Accordingly, $C_{\max}$in Algorithm 1 should be interpreted as a resource-constraint profile rather than a single computational quantity.

This high-level algorithm shows key RAFI operations: model compression to respect the budget, bias checking and mitigation, and an agentic decision step based on output confidence and context. A real implementation would refine each function (e.g., compress (), applyDebiasing(), and the agent policy π) using domain-specific methods. Crucially, at no point do we assume unlimited compute or data.

Bounded agentic operation. The RAFI agent is designed to act as an engine to support the workflow, not to completely replace clinical decision-making. Its status could include the model's prediction, uncertainty estimate, image quality data, context available in the clinical setting, device state and workflow steps which have already been completed. It can perform the following activities – request that the image be received again, ask for more information, prioritise the case for review, transfer the case to the appropriate workflow, do nothing, or escalate the case for review to a clinician. The agent cannot violate clinical safety rules.

Safety controls. RAFI needs clear procedures for its failures and escalation. When faced with problematic situations such as low quality, high uncertainty, or out-ofdistribution cases, the system should halt any activities rather than proceed with an automatic prediction. Furthermore, the deployment needs to track all the model's performance indicators, the actions taken by the agents, and the actions taken by the doctor.

3 Propositions

RAFI yields several testable propositions, which guide future evaluation studies:

- **P1:** ***Resource Efficiency vs. Deployability***. It is anticipated that selecting and refining a model based on RAFI technology will increase the number of target sites that can meet the required thresholds for latency, memory size, storage capacity, and energy consumption while maintaining acceptable diagnostic accuracy relative to the reference model.
- **P2:** ***Agentic Workflow Improves Throughput.*** In imaging workflows that require either routing or prioritisation, a limited RAFI agent is expected to deliver a faster time to review for urgent cases than passive AI assistance, as long as this doesn't result in an increase in missed urgent cases, inappropriate prioritisation, clinician override rates, and unacceptable actions of the agent.
- **P3:** ***Fairness Maintains Equity in LMIC Contexts.*** By taking fairness-related factors and subgroup monitoring into account during model development and implementation, it is assumed that RAFI would reduce gaps in clinically meaningful performance indicators among predefined demographic/clinically relevant subgroups compared to a model without fairness-related considerations. It should be emphasized that the evaluation of effectiveness in this instance has to involve such criteria as subgroup sensitivity, false-negative rate, specificity, and calibration in a necessary manner.
- **P4:** ***Continual Local Learning Yields Long-term Sustainability.*** According to RAFI, combined systems offer better results than older systems in adapting to changes in data distributions, scanners, or medical protocols, but this also does not imply that new models will start functioning immediately. The system must undergo data checks, versioning, performance testing, clinical review, and software approval before it can be implanted into the functional system. After that, a sustainability assessment can be conducted using criteria such as workshop time, performance drop, repair costs, and actual technology capacity.

The recommendations outlined above can support ongoing empirical follow-up (e.g., simulation studies or real-world pilots) without requiring fabricated data by identifying the anticipated cost savings associated with RAFI's key features and will serve as the basis for future validation of RAFI's effectiveness.

4 Results and Discussion

4.1 Comparative Conceptual Evaluation

We now contrast RAFI with conventional medical AI paradigms (centralized cloud models, standard compressed models, etc.) across several key attributes (Table 2). This is a *conceptual* comparison, not based on new experiments, but on the design analysis above.

Frugal and agentic AI may provide a practical basis for developing medical-imaging systems that are more compatible with the computational, infrastructural, and workforce constraints encountered in resource-limited healthcare settings. Frugal AI models use techniques such as compression, pruning, quantization, and edge inference to drive reductions in compute, memory, energy, and latency requirements; whereas Agentic AI is responsible for coordinating triage, providing reporting assistance, and providing clinical guidance. Together, these two types of AI technology address many common barriers faced by resource-constrained settings (e.g., lack of radiology expertise, subpar infrastructure, limited data, and unaffordable services). The integration of fairness auditing, explainability, clinical oversight, and controlled local adaptation provides mechanisms for evaluating workflow efficiency, equitable performance, and long-term maintainability. Finally, the impact of this AI framework will be significant in the long term because it provides sustainable imaging intelligence, which is crucial to improving the health of underserved communities worldwide.

Aspect	Existing AI Paradigms	RAFI Framework
Computational Efficiency	The focus of most current approaches is on developing high-quality, large models that are designed to perform well on benchmarks rather than being optimized for use in terms of runtime, power usage or deployment limitations.	RAFI uses small, lightweight models and employs dynamic inference designed to accommodate the constraints of the devices' actual hardware. To reduce FLOPs and energy consumption while maintaining diagnostic accuracy, RAFI employs techniques such as model compression, pruning, and architecture search.
Decision Autonomy	Most current AI systems are largely passive, producing outputs only when prompted, such as segmentation, classification, or detection results. Key workflow decisions still depend heavily on manual intervention.	RAFI introduces proactive agentic behaviour by automatically prioritising urgent cases, recommending protocols, and routing tasks, while maintaining clinician oversight and accountability.
Data Requirements	Conventional models generally require large, labelled datasets, often collected from well-resourced clinical settings. Their ability to adapt to new or underrepresented sites is usually limited.	RAFI promotes data-frugal learning through selective sampling, transfer learning, synthetic augmentation, and continuous local fine-tuning, allowing models to adapt effectively with limited local data.
Fairness and Equity	Bias mitigation is often treated as a post hoc concern or overlooked entirely, potentially allowing models to reproduce or amplify	RAFI integrates fairness constraints into both model design and deployment. It monitors subgroup-level performance and automatically triggers debiasing strategies when

	disparities across patient populations.	performance disparities exceed acceptable thresholds.
Deployment Model	Many existing systems are cloud-centric or rely on bulky on-premises servers, with limited attention to ultra-low-power edge deployments, such as TinyML for medical imaging.	RAFI is designed for edge or near-edge deployment using local computational resources. It prioritizes on-device inference to improve privacy, reduce latency, and enable optional cloud co-processing only for non-sensitive tasks.
Human Oversight	Human involvement is usually limited to reviewing final AI outputs. Most systems provide little support for clinician feedback, model adaptation, or local capacity building.	RAFI incorporates a human-in-the-loop structure where clinicians can override agentic decisions. It also supports educational feedback, expert-guided refinement, and local annotation of difficult or edge cases.
Sustainability	Sustainability is rarely treated as a central design objective, as most systems focus primarily on accuracy, scalability, or technical novelty.	RAFI treats sustainability as a core objective by minimizing energy use, supporting long-term model updating, tracking carbon footprints, and strengthening workforce development for global health impact.

5 Limitations and future validation.

This study presents RAFI, a newly conceived framework that has yet to be tested in clinical trials or medical imaging research. Therefore, the statements mentioned in P1-P4 can be interpreted as hypotheses only. In the future, research will validate the framework through a medical imaging task and evaluate aspects such as diagnostic capability, computing requirements, performance of subgroups, agent behavior, clinician interventions, escalation patterns, and stability after implementation. In addition, studies will investigate issues that cannot be explored conceptually, such as the influence of compression on subgroup performance and the quantifiable benefits of automating the workflow in real-life settings with limited resources.

6 Conclusion

This research presents an integrative framework (Resource-Aware Agentic Frugal Intelligence [RAFI]) to support the development of medical imaging AI in resource-constrained environments. By drawing on concepts from frugal machine learning, autonomous AI, and fairness research, RAFI addresses an urgent issue: how to make advanced analytical capabilities for illness detection using imaging data viable and accessible across diverse healthcare systems. The various design concepts, the workflow illustration (Figure 1), and the definition of abstractions within the framework constitute a rigorous and actionable blueprint. RAFI places the same amount of weight on improving efficiency, enhancing human-centered autonomy, and expanding equity; thus, from a theoretical standpoint, these objectives will generate AI systems with high clinical utility, reduced reliance on hardware, and less biased results. A key provision of the RAFI framework indicates that deployment of the technology and building capacity must be concurrent; thus, implementing this framework requires a comprehensive approach that includes training in addition to raising awareness about optimizing policy and algorithms. The merit of this paper lies in its conceptualization of a new paradigm, without providing empirical data. We invite the research community to adopt and refine RAFI's principles: to validate its propositions, to implement prototype systems, and ultimately to drive AI toward truly sustainable impact. In doing so, we can help ensure that the next generation of medical imaging AI delivers on its promise for all patients, not just those in well-resourced facilities.

References

1. Litjens, G., Kooi, T., Bejnordi, B. E., Setio, A. A. A., Ciompi, F., Ghafoorian, M., ... & Sánchez, C. I. (2017). A survey on deep learning in medical image analysis. *Medical image analysis*, *42*, 60-88.
2. Ogundokun, R. O., Misra, S., Akinrotimi, A. O., & Ogul, H. (2023). MobileNet-SVM: A lightweight deep transfer learning model to diagnose BCH scans for IoMT-based imaging sensors. *Sensors*, *23*(2), 656.
3. Strubell, E., Ganesh, A., & McCallum, A. (2019, July). Energy and policy considerations for deep learning in NLP. In *Proceedings of the 57th annual meeting of the association for computational linguistics* (pp. 3645-3650).
4. Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2020). Green ai. *Communications of the ACM*, *63*(12), 54-63.
5. Murali, S., Ding, H., Adedeji, F., Qin, C., Obungoloch, J., Asllani, I., ... & Adeleke, S. (2024). Bringing MRI to low-and middle-income countries: directions, challenges and potential solutions. *NMR in Biomedicine*, *37*(7), e4992.
6. Frija, G., Blažić, I., Frush, D. P., Hierath, M., Kawooya, M., Donoso-Bach, L., & Brkljačić, B. (2021). How to improve access to medical imaging in low-and middle-income countries?. *EClinicalMedicine*, *38*.
7. Kelly, C. J., Karthikesalingam, A., Suleyman, M., Corrado, G., & King, D. (2019). Key challenges for delivering clinical impact with artificial intelligence. *BMC medicine*, *17*(1), 195.
8. Kaissis, G. A., Makowski, M. R., Rückert, D., & Braren, R. F. (2020). Secure, privacy-preserving and federated machine learning in medical imaging. *Nature Machine Intelligence*, *2*(6), 305-311.
9. Dietrich, N. (2025). Agentic AI in radiology: emerging potential and unresolved challenges. *British Journal of Radiology*, *98*(1174), 1582-1584.
10. Faghani, S., Moassefi, M., Rouzrokh, P., Khosravi, B., & Erickson, B. J. (2025). Uncover This Tech Term: Agentic Artificial Intelligence in Radiology. *Korean Journal of Radiology*, *26*(9), 888.
11. Norori, N., Hu, Q., Aellen, F. M., Faraci, F. D., & Tzovara, A. (2021). Addressing bias in big data and AI for health care: A call for open science. *Patterns*, *2*(10).
12. Cross, J. L., Choma, M. A., & Onofrey, J. A. (2024). Bias in medical AI: implications for clinical decision-making. *PLOS digital health*, *3*(11), e0000651.
13. Larrazabal, A. J., Nieto, N., Peterson, V., Milone, D. H., & Ferrante, E. (2020). Gender imbalance in medical imaging datasets produces biased classifiers for computer-aided diagnosis. *Proceedings of the National Academy of Sciences*, *117*(23), 12592-12594.
14. Jin, R., Xu, Z., Zhong, Y., Yao, Q., Dou, Q., Zhou, S. K., & Li, X. (2024). Fairmedfm: fairness benchmarking for medical imaging foundation models. *Advances in Neural Information Processing Systems*, *37*, 111318-111357.

15. Burlina, P., Joshi, N., Paul, W., Pacheco, K. D., & Bressler, N. M. (2021). Addressing artificial intelligence bias in retinal diagnostics. *Translational Vision Science & Technology*, *10*(2), 13-13.
16. Al-Ganad, A., Al-Shahdhi, A., Al-Dhaifi, O., Hajeb, E., Hajeb, H., & Al-Motarreb, A. (2026). Deploying medical AI in low-resource settings: a scoping review of challenges and strategies. *Frontiers in Digital Health*, *8*, 1743634.
17. Dantas, P. V., Sabino da Silva Jr, W., Cordeiro, L. C., & Carvalho, C. B. (2024). A comprehensive review of model compression techniques in machine learning: PV Dantas et al. *Applied Intelligence*, *54*(22), 11804-11844.
18. Nguyen, C. M., & Hy, T. S. (2026). Efficient Deep Learning for Medical Imaging: Bridging the Gap Between High-Performance AI and Clinical Deployment. *arXiv preprint arXiv:2602.00910*.
19. Violos, J., Diamanti, K. C., Kompatsiaris, I., & Papadopoulos, S. (2026). Frugal machine learning for energy-efficient and resource-aware artificial intelligence. In *Artificial Intelligence, Data and Robotics: Foundations, Transformations and Future Directions* (pp. 175-198). Cham: Springer Nature Switzerland.
20. Lekadir, K., Frangi, A. F., Porras, A. R., Glocker, B., Cintas, C., Langlotz, C. P., ... & Starmans, M. P. (2025). FUTURE-AI: international consensus guideline for trustworthy and deployable artificial intelligence in healthcare. *bmj*, *388*.
21. Collaco, B. G., Haider, S. A., Prabha, S., Gomez-Cabello, C. A., Genovese, A., Wood, N. G., ... & Forte, A. J. (2026). The role of agentic artificial intelligence in healthcare: a scoping review. *npj Digital Medicine*.
22. Koçak, B., Ponsiglione, A., Stanzione, A., Bluethgen, C., Santinha, J., Ugga, L., ... & Cuocolo, R. (2025). Bias in artificial intelligence for medical imaging: fundamentals, detection, avoidance, mitigation, challenges, ethics, and prospects. *Diagnostic and interventional radiology*, *31*(2), 75.
23. World Health Organization. (2024). Ethics and governance of artificial intelligence for health: large multi-modal models. WHO guidance. World Health Organization.
24. Unesco. (2022). *Recommendation on the ethics of artificial intelligence*. United Nations Educational, Scientific and Cultural Organization.
25. Mollura, D. J., Culp, M. P., Pollack, E., Battino, G., Scheel, J. R., Mango, V. L., ... & Dako, F. (2020). Artificial intelligence in low-and middle-income countries: innovating global health radiology. *Radiology*, *297*(3), 513-520.