



Foundation models for radiology: fundamentals, applications, opportunities, challenges, risks, and prospects

Tugba Akinci D'Antonoli^{1,2}

Christian Bluethgen³

Renato Cuocolo⁴

Michail E. Klontzas^{5,6,7}

Andrea Ponsiglione⁸

Burak Kocak⁹

¹University Hospital Basel, Department of Diagnostic and Interventional Neuroradiology, Basel, Switzerland

²University Children's Hospital Basel, Department of Pediatric Radiology, Basel, Switzerland

³University Hospital Zurich, Institute for Diagnostic and Interventional Radiology, Zurich, Switzerland

⁴Department of Medicine, Surgery and Dentistry, University of Salerno, Baronissi, Italy

⁵University of Crete School of Medicine, Artificial Intelligence and Translational Imaging (ATI) Lab, Department of Radiology, Crete, Greece

⁶Foundation for Research and Technology (ICS-FORTH), Institute of Computer Science, Computational Biomedicine Lab, Crete, Greece

⁷Karolinska Institute, Department of Clinical Science, Intervention and Technology (CLINTEC), Division of Radiology, Stockholm, Sweden

⁸University of Naples Federico II, Department of Advanced Biomedical Sciences, Naples, Italy

⁹Başakşehir Çam and Sakura City Hospital, Department of Radiology, İstanbul, Türkiye

Corresponding author: Tugba Akinci D'Antonoli

E-mail: tugba.akincidantonoli@unibas.ch

Received 06 May 2025; revision requested 10 June 2025; accepted 22 June 2025.



Epub: 08.07.2025

Publication date: xx.07.2025

DOI: 10.4274/dir.2025.253445

ABSTRACT

Foundation models (FMs) represent a significant evolution in artificial intelligence (AI), impacting diverse fields. Within radiology, this evolution offers greater adaptability, multimodal integration, and improved generalizability compared with traditional narrow AI. Utilizing large-scale pre-training and efficient fine-tuning, FMs can support diverse applications, including image interpretation, report generation, integrative diagnostics combining imaging with clinical/laboratory data, and synthetic data creation, holding significant promise for advancements in precision medicine. However, clinical translation of FMs faces several substantial challenges. Key concerns include the inherent opacity of model decision-making processes, environmental and social sustainability issues, risks to data privacy, complex ethical considerations, such as bias and fairness, and navigating the uncertainty of regulatory frameworks. Moreover, rigorous validation is essential to address inherent stochasticity and the risk of hallucination. This international collaborative effort provides a comprehensive overview of the fundamentals, applications, opportunities, challenges, and prospects of FMs, aiming to guide their responsible and effective adoption in radiology and healthcare.

KEYWORDS

Artificial intelligence, deep learning, foundation models, ChatGPT, large language models, medical imaging, radiology

Artificial intelligence (AI), particularly deep learning (DL), has demonstrated considerable efficacy in medical image analysis across various imaging modalities.^{1,2} Traditionally, however, AI models in healthcare have been mostly developed for narrow tasks that are highly specific and limited. The recent emergence of foundation models (FMs) represents a significant paradigm shift.^{3,4} These large DL models exhibit broad adaptability to a wide range of downstream tasks with minimal task-specific modification.^{5,6}

A notable example of FMs is represented by large language models (LLMs), optimized for language-centric tasks, such as summarization, translation, and answering questions.⁷ Although LLMs primarily process text, the broader category of FMs can encompass multiple modalities, including text, images, audio, and a diverse spectrum of unstructured data.^{8,9} This inherent multimodality aligns well with the diverse data types encountered in modern medicine, such as imaging, clinical narratives, laboratory results, and genomic information.¹⁰⁻¹²

Although current radiology workflows predominantly utilize task-specific models, the multimodal capabilities of FMs make them particularly promising for this field, offering potential support across various interpretative and non-interpretative scenarios (Figure 1).^{13,14} The capabilities of LLMs have already been explored for several radiology-related tasks, including report generation,¹⁵ multilingual report translation,¹⁶ information extraction from free-text reports,¹⁷ and the assessment of domain-specific radiological knowledge.¹⁸ Despite growing interest, the use of FMs in radiology is still in the early stages, with ongoing active research and development.^{13,14,19-21}

To facilitate the development and potential adoption of FMs, this narrative review synthesizes current knowledge about FMs and aims to provide a comprehensive overview of FMs in the context of radiology. It introduces the fundamental concepts behind FMs, examines their potential applications in radiology, highlights emerging opportunities, outlines key challenges, and suggests future directions in both research and practice.

Fundamental concepts of foundation models

FMs mark a fundamental shift within the conceptual hierarchy of AI (Figure 2), moving beyond conventional, narrowly focused AI systems. They are a class of large-scale AI models developed through training on vast and diverse datasets.³ A defining feature is their pre-trained nature; unlike conventional models engineered for a single, narrow task (e.g., solely lung nodule detection or lung segmentation), FMs serve as versatile base models (created through pre-training), adaptable to numerous downstream applications (through fine-tuning, continued training on smaller, task-specific datasets).

This inherent adaptability results from several key characteristics (Figure 3). First, the pre-training stage of FMs usually leverages self-supervised learning, allowing the model to learn rich data representations from the data itself (e.g., by solving pretext tasks, such as predicting masked portions of an image or a text) using unstructured, unlabeled, or weakly labeled data (Figure 4).^{22,23} This contrasts sharply with conventional methods, which typically require significant amounts of high-quality (manually) labeled data for each distinct task, a major bottle-

neck due to cost and expert time. Although FMs still require labeled data for fine-tuning, the reliance on specific data for each application can be substantially reduced compared with conventional AI methods (Figure 5).

Second, the reduced need for labeled data allows FM development on a large scale, referring to model size, computational resources, and dataset size.²⁴ This scale enables the models to learn more generalizable and robust representations that, in turn, support the scalability of the FMs themselves (larger models become feasible with more training data, and more generalizable representations apply to more possible downstream applications). This differentiates FMs from conventional models, which often exhibit limited generalizability beyond the precise conditions (e.g., patient populations or tasks) for which they were trained. Additionally, scaling models has led to the emergence of functionalities beyond the explicit training objectives,³ such as instruction following, capabilities unprecedented (or at least hardly detectable) in smaller-scale models.²⁵⁻²⁷

Finally, self-supervised learning and the scale of FMs equip them with strong transfer learning capabilities.²⁸⁻³¹ The general knowledge acquired during the resource-intensive pre-training phase can be effectively utilized for new, specific tasks through minimal fine-tuning. This facilitates few-shot learning (where only a small number of task-specific examples are provided) and zero-shot learning (using no examples),³²⁻³⁴ where models adapt with substantially less specific data than conventional approaches demand. For instance, an FM pre-trained via self-supervised learning on large chest X-ray datasets may be fine-tuned for rib fracture detection using only dozens of cases, whereas a conventional model may require thousands to reach comparable performance.

Developing multimodal foundation models

FMs first took shape in natural language processing (NLP) in the form of LLMs, such as Generative Pre-trained Transformer (GPT)-4 (OpenAI) and Claude (Anthropic). Although FMs can be unimodal, focusing exclusively on one data type, such as text (in the case of LLMs) or images,^{35,36} a development of particular importance for radiology is their potential to be multimodal by being able to process and integrate diverse data types, including images [e.g., X-rays,³⁷⁻⁴⁰ computed tomography (CT),^{41,42} and magnetic resonance imaging (MRI)⁴³], text (e.g., reports and

other electronic health record documents), and potentially many more (Figure 1).

Key concepts and modules of FMs concerning radiological applications are presented in Figure 6. Although architectures vary, the transformer design is a frequently used backbone.⁴⁴ Its central feature, the attention mechanism, allows it to focus on specific elements of the input sequence. This enables the model to capture long-range dependencies and contextual relationships within data effectively, which gave rise to its initial success in NLP and subsequent adaptation for vision and multimodal scenarios.^{45,46} A key concept in handling diverse inputs is the use of modality-specific encoders.⁴⁷ These components compress high-dimensional inputs (such as CT scans or text reports) into lower-dimensional embeddings (i.e., numerical vector representations), capturing essential features (e.g., tissue density, anatomical structures, radiological terms). Common encoder architectures include vision transformers and convolutional neural networks for images, and transformers for processing text data. To enable the model to understand relationships across different data types, techniques such as contrastive learning are often employed during pre-training (Figure 7). For instance, the model learns that a specific chest X-ray and its corresponding report describe the same case. Model weights are adjusted so that the embeddings for a matching image-report pair are pulled closer together in a conceptual "shared space," whereas embeddings for unrelated pairs (e.g., the same chest X-ray paired with a report from a different patient) are pushed further apart to learn meaningful cross-modal associations. The Contrastive Language-Image Pre-Training (CLIP) model is a dual neural network trained on a variety of image and text data pairs and is an early example of FMs created this way.⁴⁵

After the individual encoders have processed their respective inputs, fusion modules are used to combine this information, which can happen in several ways (Figure 8). Mechanisms such as cross-attention are particularly powerful here, allowing the model to weigh dynamically the relevance of different parts of one modality based on the content of another—for example, attending to specific words in a report when analyzing a corresponding slice in a CT scan.

Finally, decoders transform these fused representations into desired outputs, which could range from generating text (e.g., report summaries) and predicting classes or outcomes to segmenting relevant image

Main points

- Foundation models (FMs) are versatile artificial intelligence (AI) systems pre-trained on large, diverse datasets, enabling them to adapt to many tasks with minimal fine-tuning.
- FMs with multimodal capacities offer powerful tools for complex radiological applications, such as report generation and diagnostic decision-making.
- FMs have the potential to democratize AI in healthcare by requiring less local data for fine-tuning, helping under-resourced centers.
- Major challenges to FM use in imaging include stochasticity, hallucinated outputs, transparency, bias, sustainability, and regulations.

Foundation Models in Radiology

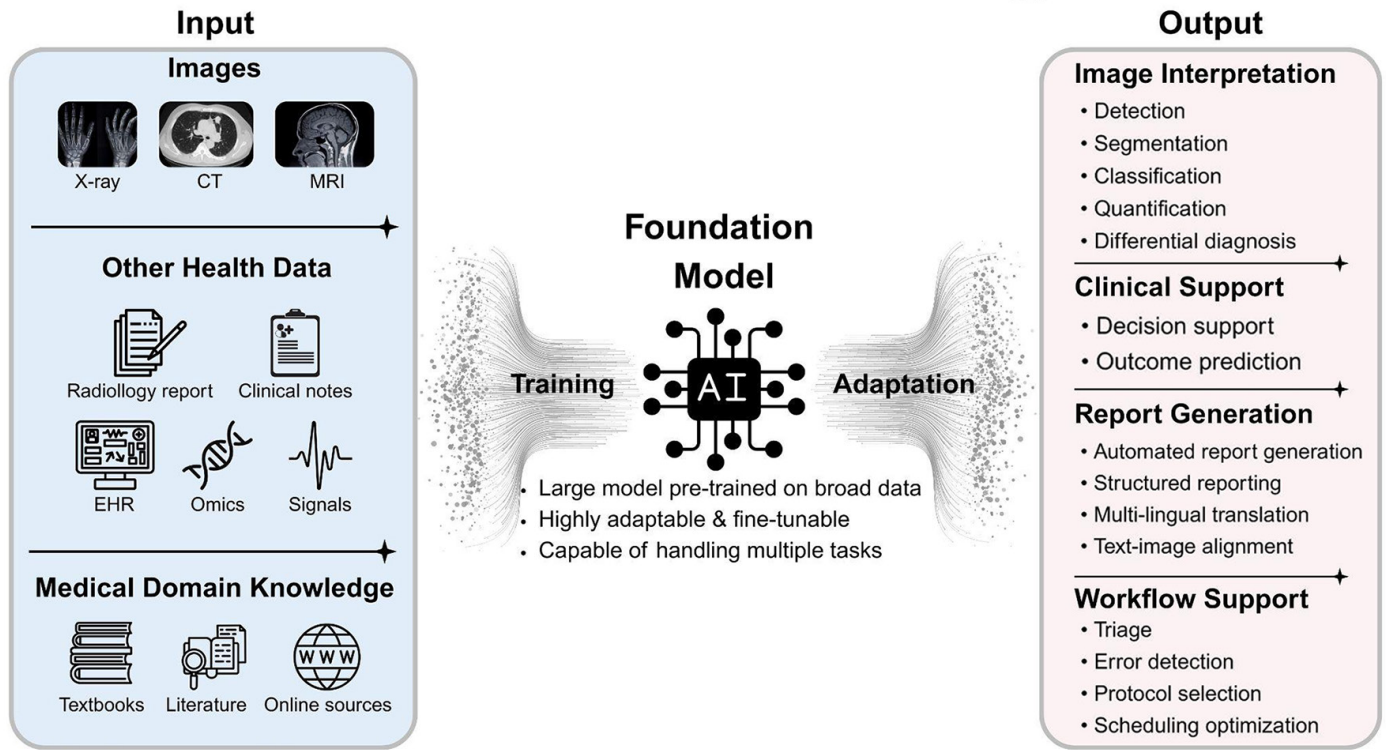


Figure 1. Foundation models in radiology with diverse inputs to perform various interpretative and non-interpretative tasks through training and adaptation. CT, computed tomography; MRI, magnetic resonance imaging; AI, artificial intelligence; EHR, electronic health record.

regions, depending on the specific application. Adapting existing FMs to a specific task can be achieved through full fine-tuning (updating all model parameters) or less computationally expensive parameter-efficient fine-tuning techniques that update only a small number of parameters (Figure 9).⁴⁸

These characteristics make FMs versatile, adaptable, and data-efficient AI models that can integrate multimodal data and capture long-range dependencies within high-dimensional data and across different modalities that may elude narrower models. This uniquely positions FMs to tackle complex problems from the field of radiology by providing potentially richer, more contextualized insights that reflect clinical reality better than conventional AI models.⁴⁷

Applications in radiology

Creating a radiology-specific FM from scratch could be highly cost-intensive, as radiology consists of a wide range of imaging modalities, including X-rays, ultrasound, nuclear imaging, and MRI, that have significant variations in their underlying technologies and data characteristics.⁴⁹ Nevertheless, recent advances have shown promising pathways to adapt or fine-tune general-purpose models for domain-specific tasks and mo-





Figure 2. Oversimplified conceptual hierarchy, illustrating the relationship and progression of key domains in artificial intelligence towards foundation models with multimodal capabilities. AI, artificial intelligence.

dalities, leading to a growing number of successful applications in radiology.

One core application is medical image segmentation, which aims to delineate re-

gions of interest, such as lesions or organs, automatically. DL models, especially those using the nnU-Net architecture, have shown high accuracy in normal anatomical seg-

Key Characteristics of Foundation Models




Scalability & Emergence

- Trained at massive scale enabled by low-label data
- Generalizes across domains with robust pattern recognition
- Exhibits emergent abilities not explicitly trained (e.g., instruction following)

Self-Supervised Learning

- Learns from raw data via pretext tasks
- Minimizes dependence on manual labeling
- Reduces data annotation costs
- Accelerates model development



Strong Transfer Learning

- Applies pre-trained knowledge to new tasks
- Requires minimal fine-tuning
- Adapts with few examples (i.e., few-shot learning)
- Performs tasks without any examples (i.e., zero-shot learning)
- Superior data efficiency (i.e., less task specific data) over traditional models




Figure 3. Summary of three key characteristics of foundation models.

Self-Supervised Learning Examples

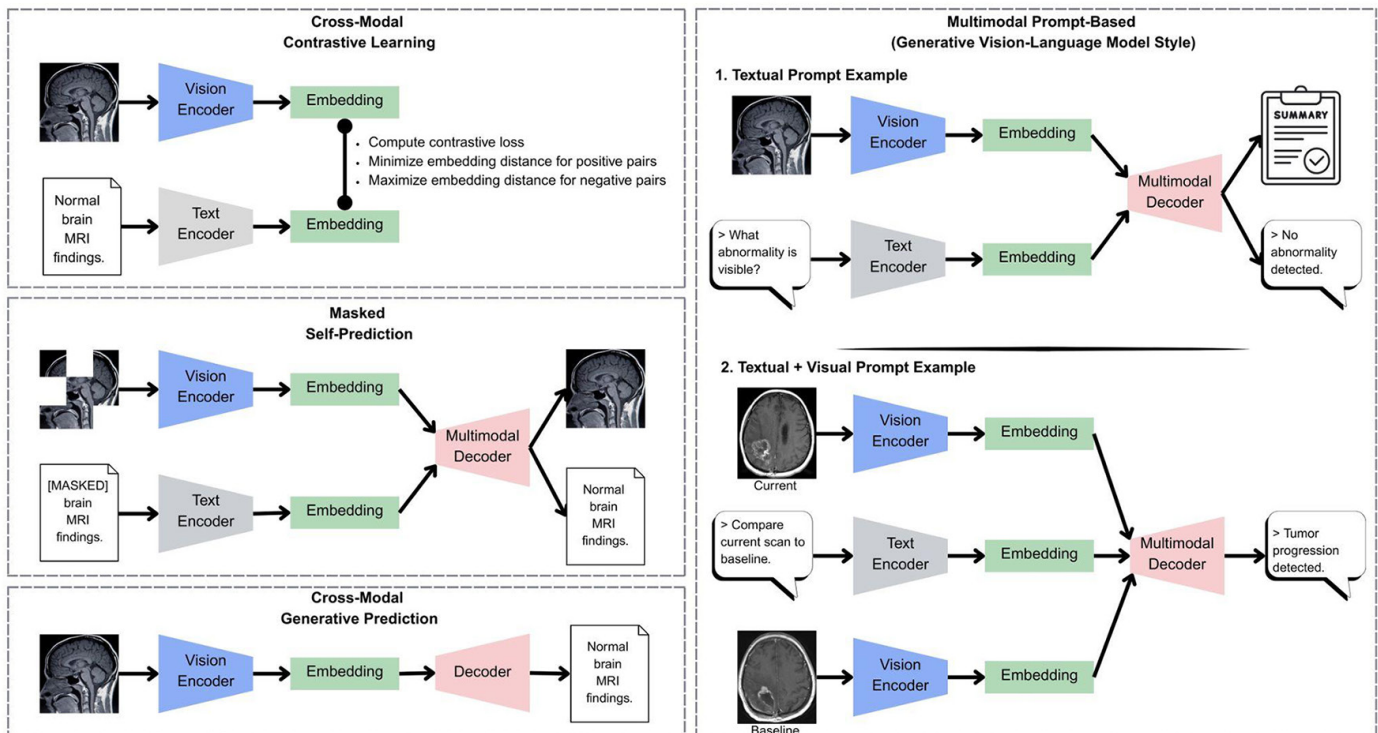


Figure 4. Simplified multimodal self-supervised learning examples. (Top left) Contrastive learning aligns paired image-text embeddings [e.g., brain magnetic resonance imaging (MRI) and report], usually by minimizing the distance between positive pairs and maximizing the distance between negative pairs. (Middle left) Masked self-prediction reconstructs masked regions of an image or text using cross-modal context. (Bottom left) Cross-modal generation predicts one modality from another (e.g., full text reports from image embeddings). (Top right) Text-only prompt-based learning uses a natural language query (e.g., “What abnormality is visible?”) to guide generative outputs based on an MRI image input. (Bottom right) Multimodal prompt-based learning uses both current and baseline imaging along with text prompts (e.g., “Compare current scan to baseline”) to support complex clinical tasks, such as disease progression assessment or change detection. The baseline image serves as a visual prompt.

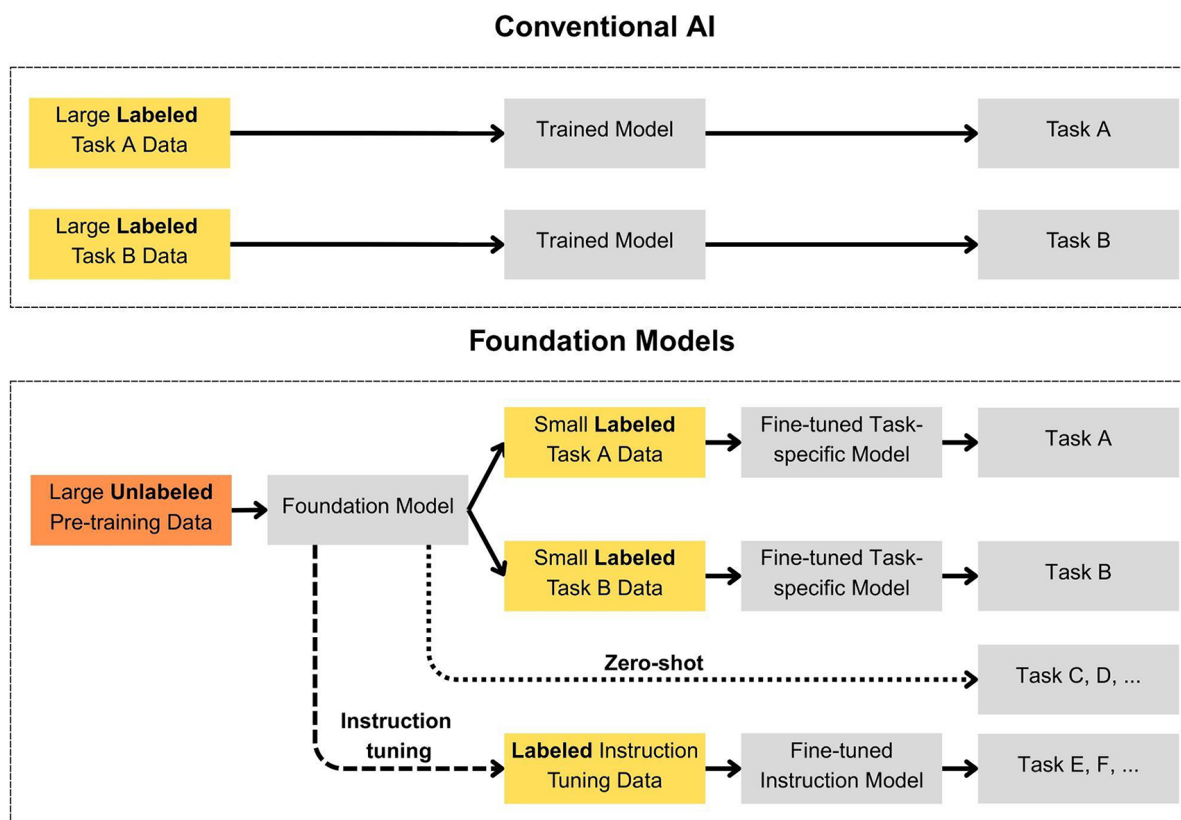


Figure 5. Comparison of conventional artificial intelligence (AI) and foundation models (FMs). Conventional AI requires large, labeled datasets tailored to each specific task and produces models limited to that task. For example, Task A (e.g., lung nodule detection on chest X-ray) and Task B (e.g., lung opacity classification) each require separate models trained on task-specific labeled data. In contrast, FMs are pre-trained on massive unlabeled data and can be adapted to various downstream tasks using smaller labeled datasets. FMs also support zero-shot learning, where the model performs new tasks, such as Task C (e.g., pneumonia detection) and Task D (e.g., pleural effusion identification) without fine-tuning or labeled data. Instruction tuning allows the model to learn from instruction–output pairs (i.e., prompt > output or prompt + image > output) and match expected outputs by adjusting itself when errors occur. This enables it to perform complex tasks, such as Task E (e.g., automated lung report generation) and Task F (e.g., triage of abnormal chest findings) without needing labeled examples for each.

mentation tasks.^{50–52} However, the challenge arises when it comes to pathologies, where separate models need to be trained for each one of them; for example, a model trained for liver tumor segmentation cannot be directly applied to the lung tumor or prostate cancer segmentation tasks.

In April 2023, the Segment Anything Model (SAM) was introduced, demonstrating the potential for a single model to handle various segmentation tasks across different domains without needing retraining or fine-tuning.⁵³ Despite this, SAM’s performance on complex medical segmentation tasks, such as those involving the pancreas and spine, has not been satisfactory.^{51,54,55} Vision FMs, such as SAM, can serve as starting points, which are then adapted into modality/task-specific FMs (e.g., Medical SAM 2) optimized by leveraging the unique characteristics of medical imaging modalities.^{37,56,57}

Another notable development is UniSeg, a single task-agnostic model trained using a large and diverse set of open-access medical datasets.⁵⁸ This model can general-

ize to the anatomies and segmentation tasks that were not in the training set or were never seen by the model previously. Notably, UniSeg significantly outperformed existing few-shot methods across all held-out datasets. However, it is important to note that the model was only applied to two-dimensional data and single-label segmentation, and its performance for three-dimensional medical image data remains unclear.

FMs have also shown promise for lesion identification and characterization in different clinical scenarios.^{19,59} For example, CXR-Base was developed using a large collection of unlabeled chest X-ray images through self-supervised learning.⁵⁹ This approach was sequentially applied to both natural images from ImageNet-1k and chest X-ray images from various public datasets, encompassing a total of 1.1 million chest X-ray images. The model demonstrated good performance across multiple datasets from different centers for diagnosing diseases such as coronavirus disease 2019, pneumonia, and tuberculosis.

Radiology report generation and comprehension represent further promising areas for multimodal FMs. In the task of generating radiology reports, these models can identify abnormalities within images from various modalities while incorporating the patient’s medical history and clinical examination findings.⁶⁰ By integrating both text and images, these models can generate precise radiology reports, help standardize report quality by detecting inconsistencies or omissions, and subsequently reduce the workload for radiologists.⁶¹ Additionally, these models can generate reports in multiple languages and adjust the complexity of the language to suit the target audience, providing detailed content for specialists and simplified versions for general practitioners.^{16,62}

For the comprehension task, physicians can also use multimodal models to enhance case comprehension by engaging in text-based dialogues that focus on specific image sections, allowing for detailed descriptions of those areas.⁶³ Furthermore, the reports generated by FMs can offer preliminary diag-

Key Concepts and Modules of Foundation Model Architecture for Radiology AI

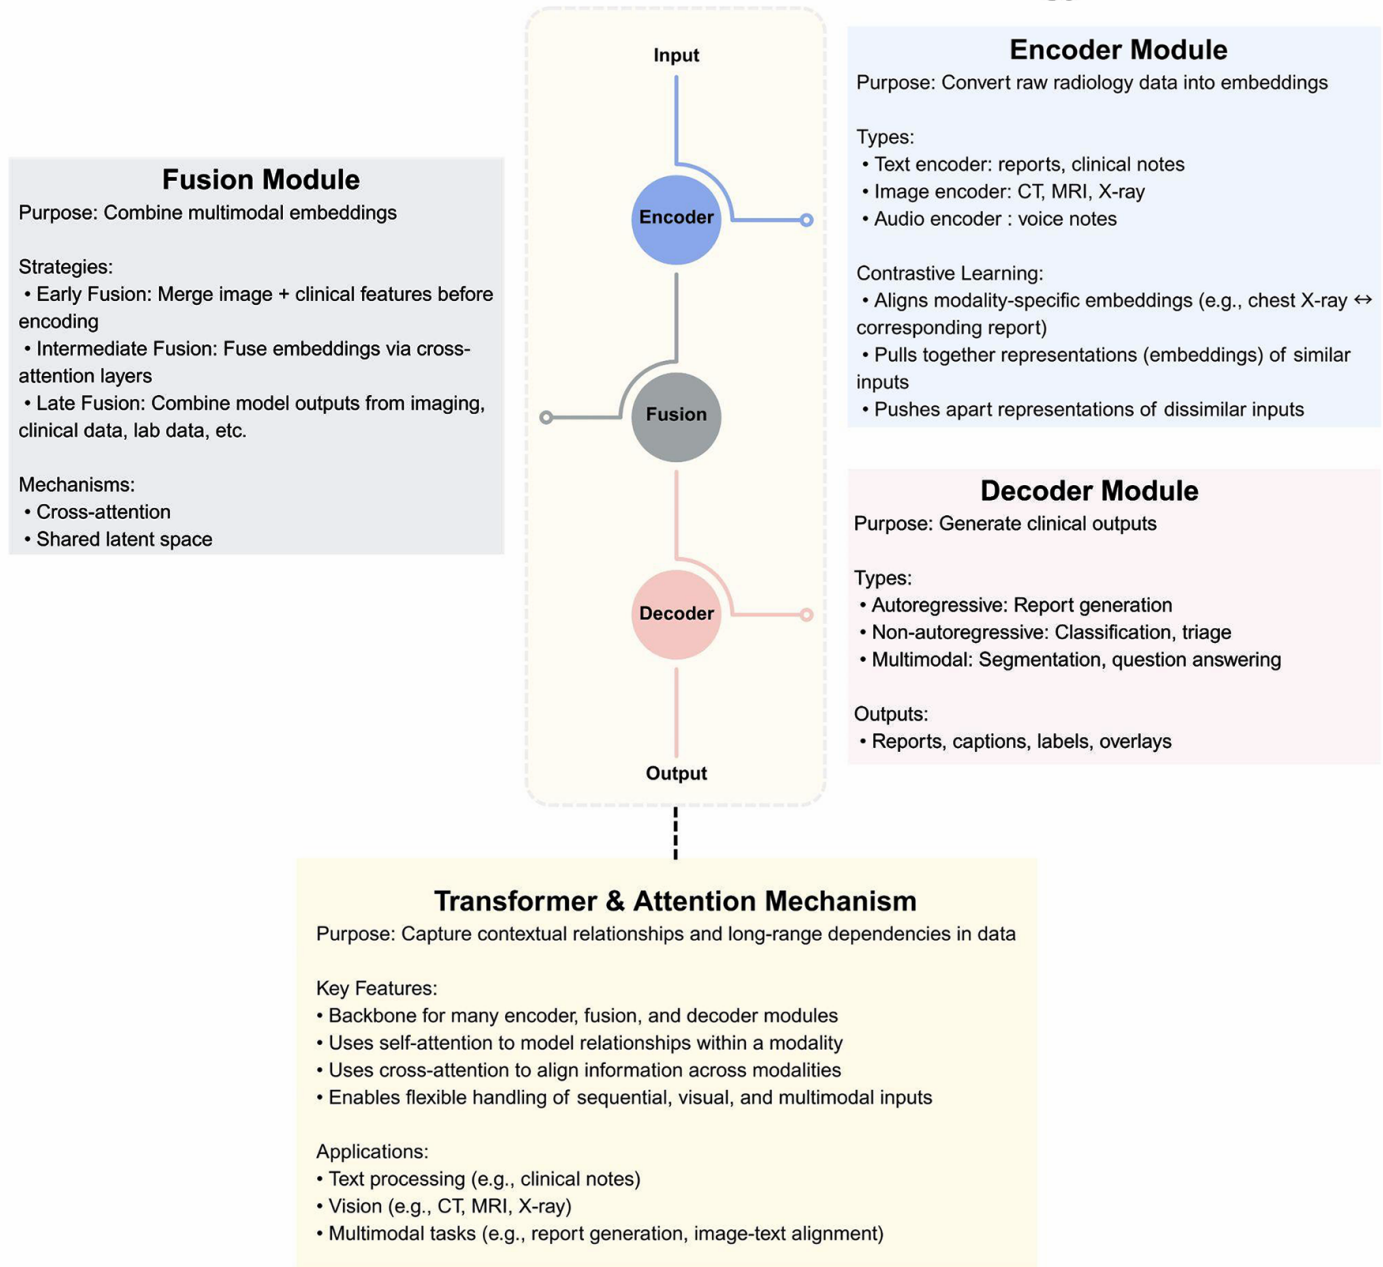


Figure 6. Key concepts and modules of foundation models, with respect to radiological applications. CT, computed tomography; MRI, magnetic resonance imaging; AI, artificial intelligence.

noses, supporting radiologists and clinicians in their decision-making process.⁶⁴ These models can potentially propose treatment options or recommend additional diagnostic tests, enhancing the overall clinical workflow for personalized medicine.⁶⁵

Opportunities for radiology

Building on the capabilities outlined above, the adoption of FMs in radiology presents several strategic opportunities (Figure 10).

A key opportunity lies in the ability to fine-tune pre-trained FMs with smaller, local datasets.⁶⁶ This can reduce the inequalities related to the availability of data,⁶⁷ democratizing access to these applications for health-care systems with limited data or access to infrastructure. The lower reliance on large datasets can allow their use in centers with limited funding or limited population coverage, which prevents the collection of a high number of cases.

Fine-tuning the models with local datasets can mitigate biases related to underrep-

resented population characteristics or local peculiarities related to equipment or radiological protocols. At the same time, leveraging their pre-training on large datasets, FMs trained on diverse populations can provide more equitable care recommendations, reducing diagnostic errors in underrepresented groups such as children, ethnic minorities, or patients with rare conditions.⁶⁸

To further address data imbalance, techniques such as synthetic data generation can be used. FMs have the potential to create synthetic medical images, such as CT scans,

Contrastive Learning

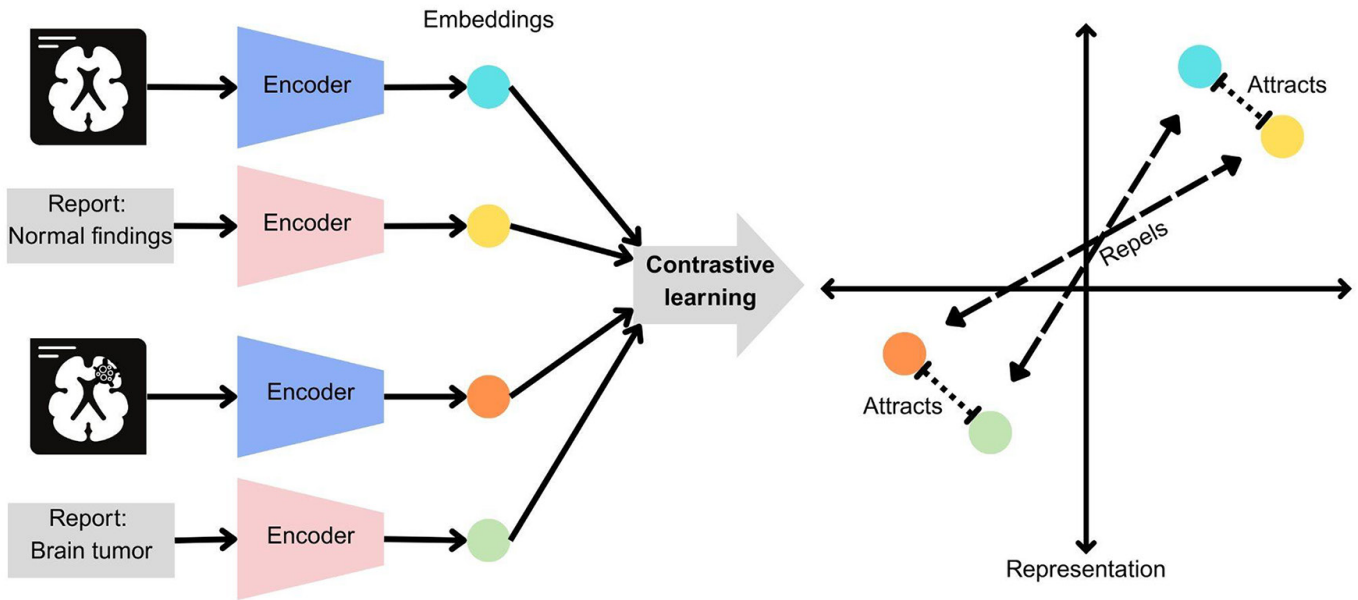


Figure 7. Simplified illustration of contrastive learning. Paired medical images and reports are encoded into embeddings (i.e., numerical vector representations), which are trained to attract matching pairs and repel mismatched ones. This process aligns similar representations in the shared embedding space to improve cross-modal understanding.

MRIs, and X-rays, that resemble real-world data.^{69,70} These artificially generated datasets can serve as valuable supplements to existing image collections, particularly when access to patient data is restricted due to privacy issues or limited availability. By generating variations of medical images, these models can help address imbalances in datasets, effectively representing a broader spectrum of pathologies.

The ability to build upon pre-trained backbones has the potential to shorten the innovation-to-implementation cycle significantly. Researchers can build on top of models trained to capture broad medical imaging features and clinical context, rather than creating new models from scratch.⁷¹ Reducing the duration of the innovation-to-implementation cycle can accelerate the development of novel applications, simplify cross-institutional collaborations, and allow innovations developed in academic settings to be rapidly tested and adapted in hospitals, startups, or public health agencies.

FMs also offer unique educational benefits. Automatically annotating synthetic or real images with detailed descriptions, such as the identification of lesions, tumors, or anatomical landmarks, can help radiology residents and healthcare professionals quickly understand complex images. Furthermore, FMs can enhance educational content by not only annotating images but also offering

Fusion Types

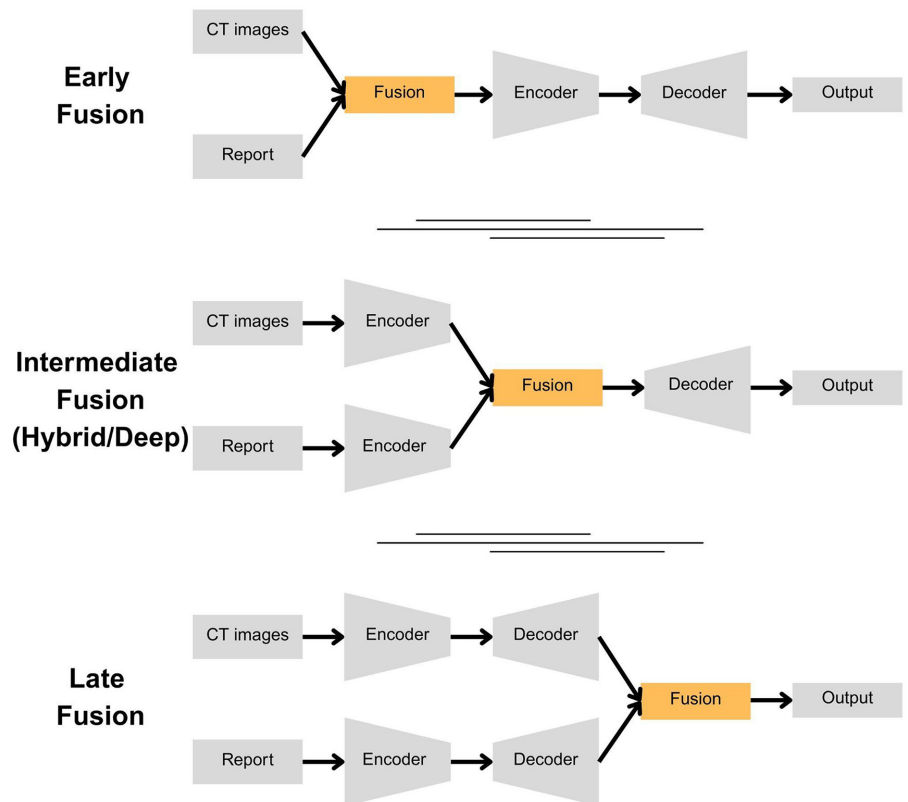


Figure 8. Simplified illustration of multimodal fusion strategies. Illustration of early, intermediate, and late fusion methods for combining medical images and their respective reports. Intermediate fusion (i.e., hybrid or deep fusion) is more common in foundation models due to its flexibility and effectiveness in capturing cross-modal interactions. CT, computed tomography.

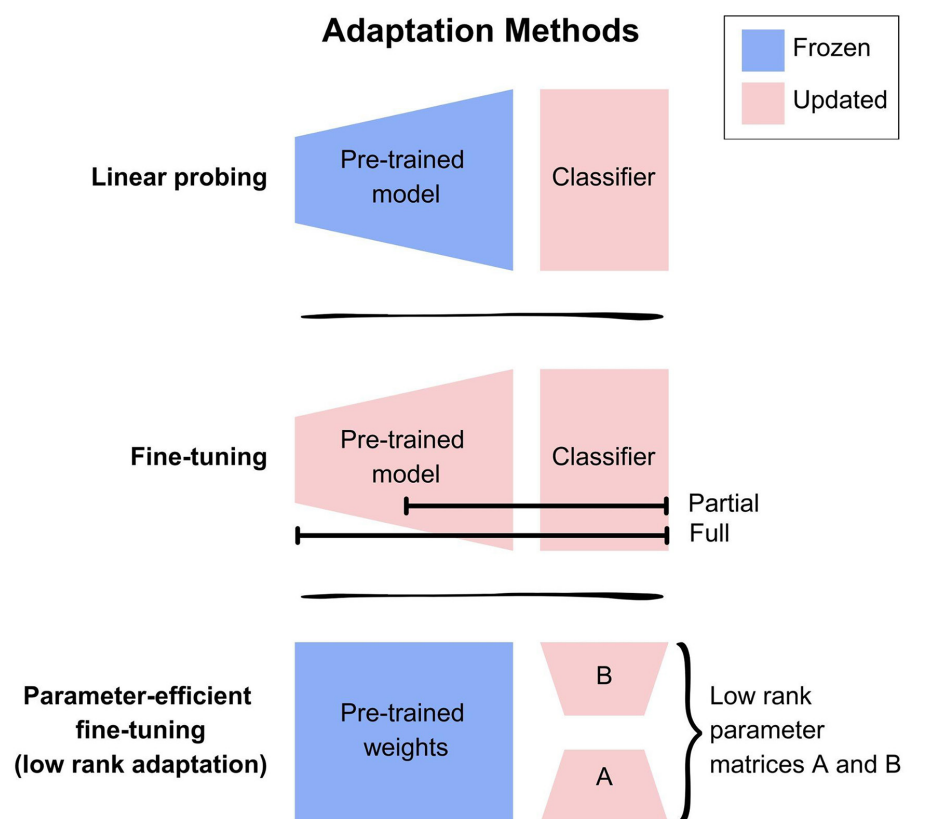


Figure 9. Simplified overview of adaptation methods for pre-trained models. Linear probing updates only the classifier, fine-tuning updates the classifier and optionally parts or all of the pre-trained model, whereas parameter-efficient fine-tuning (e.g., LoRA) inserts low-rank matrices (A and B) without updating original weights.

their condition, treatment options, and possible outcomes with FMs, thus helping patients gain a clearer understanding of their specific health situation.^{72,73}

FMs introduce a new paradigm in precision medicine by enabling the integration of diverse data types, such as the combination of diagnostic images, omics, clinical, and laboratory data.¹⁰ Radiologists interpret images in light of clinical information and questions and combine or harmonize these different types of information without much effort in their day-to-day jobs. However, existing AI applications have been inherently less accurate than radiologists, wherever such data harmonization was required. With the advance of FMs, there has been a significant leap towards the combination of multimodal data, which enables more accurate prognostication, risk stratification, and treatment planning.⁷⁴

Beyond data integration, FMs also have the capacity to support complex diagnostic reasoning in uncertain or ambiguous clinical situations. In real clinical practice, radiologists often deal with cases that are not clear-cut, where the diagnosis is not obvious, and decisions have to be made with incomplete information, taking into account follow-up data. FMs can execute different tasks that can help in these situations by offering insights that take the full clinical context into account, rather than just giving a simple yes/no answer.⁴⁷ Such models can be trained on a variety of tasks representing real diagnostic scenarios, where the radiologist is presented with imaging examinations that depict a series of pathological conditions and require complex reasoning.

To provide a perspective on emerging opportunities, a comparative overview of conventional AI (single or multimodal), multimodal FMs, radiologist interpretation, and a combined radiologist-FM approach is presented in Table 1.

Challenges and risks for radiology

As discussed so far, although FMs hold promise for transforming radiology, they also introduce multifaceted challenges. Radiologists must remain aware of these issues and proactively address them to ensure the safe, ethical, and effective implementation of FMs in radiology (Figure 11).

One of the main concerns is the stochastic nature of these models, where their outputs may vary every time they are executed, and another is that they can generate plausi-

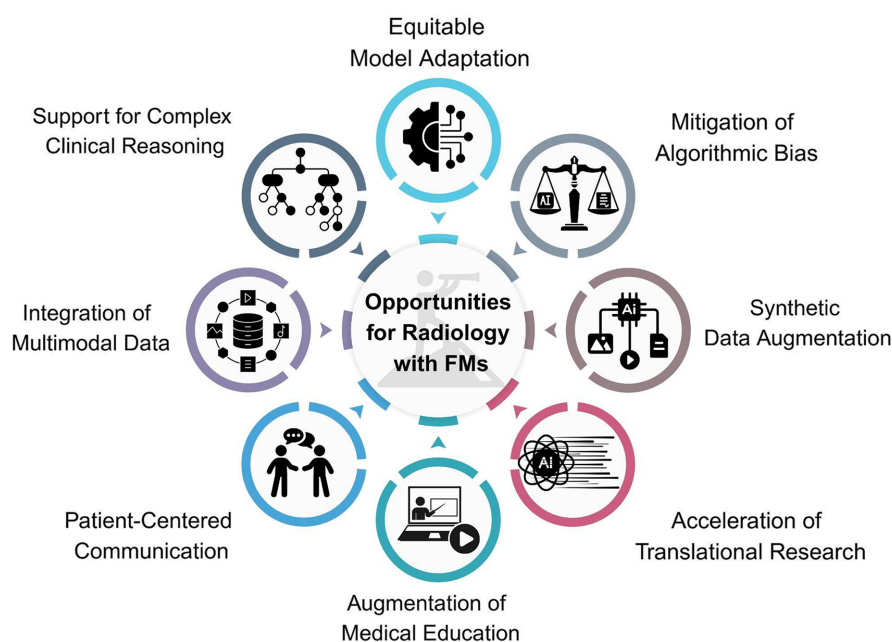


Figure 10. Potential opportunities of foundation models in radiology. FM, foundation models.

detailed explanations of pathologies, their clinical importance, and treatment options, creating an engaging and interactive learning experience.²¹

Moreover, patients will be empowered as their health data, such as their medical history and imaging results, could be translated into personalized, simplified explanations of

Table 1. Comparative overview of hypothetical glioma diagnostic pathways using conventional artificial intelligence, multimodal foundation models, radiologist interpretation, and a combined radiologist–FM approach, with an optimistic perspective

Aspect	Conventional AI (single or multimodal)	Multimodal FM	Radiologist (human)	Radiologist + multimodal FM
Inputs	Brain magnetic resonance imaging only + structured clinical data (e.g., age, simple symptoms)	MRI + clinical notes (e.g., seizures, cognitive decline) + structured labs (e.g., IDH status if known) + prior imaging	MRI + full clinical history + prior scans + surgical reports, pathology if available	Same as FM (MRI + notes + labs + prior imaging)
Processing	Task-specific detection/classification (e.g., identification of mass, segmentation of lesion, basic grading)	Integration of MRI patterns + textual symptoms + past treatments to predict diagnosis, with multimodal reasoning	Correlation of imaging patterns with clinical context and history (e.g., tumor vs. abscess vs. gliosis)	FM preliminary synthesis + human critical review, correction, and clinical judgment
Clinical understanding	Limited contextual understanding detects and segments tumor but lacks the ability to interpret complex clinical nuances not trained for	Broader contextual reasoning by combining multiple modalities, such as text and image; may suggest differential diagnoses, but reliability varies, particularly in rare or atypical cases	Deep contextual understanding built from experience and training; able to integrate subtle imaging findings with complex clinical backgrounds	Combines FM rapid synthesis with human clinical judgment; radiologist can verify, contextualize, and override FM suggestions when needed
Decision support	Reports lesion size, volume, location, edema; sometimes estimates basic tumor type (low-grade vs. high-grade)	Provides full contextual report (e.g., type prediction, treatment suggestions such as re-operation, re-irradiation)	Provides full diagnostic impression (e.g., tumor type suggestion, progression vs. pseudoprogression analysis, biopsy guidance)	FM provides draft; radiologist validates, corrects, adds nuanced interpretation, finally providing an optimized report
Adaptability	Low to moderate: Cannot dynamically adapt to unexpected cases (e.g., rare glioma subtypes, atypical presentations)	High: Generalizes from large data, but can still hallucinate or miss rare conditions	Very high: Experienced ones adapt to rare, complex, or conflicting clinical/imaging situations	Potentially highest: However, real world performance depends on the quality of human–AI interaction (e.g., hesitation to override AI outputs)
Interaction	No interaction: Static outputs	Emerging: Can answer prompts (e.g., “What is differential diagnosis?”)	High: interacts with clinicians, neurosurgeons, oncologists, and pathologists; discusses complex cases	Highest: Radiologist queries FM, adjusts outputs, discusses with clinicians dynamically
Limitations	Misses context (e.g., prior radiation therapy effects); cannot explain atypical findings	Hallucination risk; lack of explainability in complex or rare cases	Human variability, fatigue, long interpretation times	Human-in-the-loop improves safety, but challenges remain (e.g., bias, liability concerns, need for clear definitions of radiologists’ role in AI-assisted decision support)

AI, artificial intelligence; FM, foundation models; MRI, magnetic resonance imaging; IDH, isocitrate dehydrogenase.

ble-sounding yet incorrect or entirely fabricated information (a phenomenon known as “hallucination”).⁷ Beyond these inherent issues, FMs present several broader challenges, especially in radiology and healthcare in general. These include challenges related to sustainability, transparency, ethics, cybersecurity, privacy, standardization, and validation.

Another concern is sustainability and environmental impact, as the development and deployment of FMs are highly resource-intensive. They demand vast computational power, energy, and even water. For example, generating a single image with a generative AI model can consume the equivalent of half a smartphone’s battery charge,⁷⁵ while producing 10–50 medium-length chatbot responses may require up to half a liter of fresh water (Figure 12).⁷⁶ Since radiology already depends on energy-intensive imaging equipment, implementing these models in

radiology could further exacerbate environmental burdens.^{77,78} Addressing these issues and promoting sustainable practices are essential to reducing the environmental impact of FMs in radiology.

FMs often operate as “black boxes,” producing outputs without providing clear explanations of their reasoning processes.⁷⁹ In radiology, transparency and explainability are critical. Diagnostic decisions made by radiologists must be evidence-based to guide treatment plans and ensure patient safety. If a radiologist cannot justify a diagnosis aided by these tools, their trust in this decision may be undermined. Implementing models with reasoning capabilities, such as OpenAI’s GPT o1 series, or adopting frameworks that are designed to facilitate reasoning may help improve trust in model outputs.⁸⁰

Moreover, FMs are prone to perpetuating or even amplifying biases present in the training data, which can contribute to

healthcare disparities.⁶⁷ In addition, unequal access to such technologies may further disadvantage under-resourced institutions.⁶⁷ The legal framework surrounding AI in radiology is also still evolving, and questions about liability, especially in cases where diagnostic errors result from following or ignoring AI-aided recommendations, remain unresolved.⁸¹ Clear ethical guidelines and legal standards are needed to navigate these challenges responsibly.

Training FMs requires large datasets that may contain sensitive patient information. This raises substantial privacy concerns; therefore, ensuring rigorous data anonymization practices during the model training, as well as not using patient data directly as input during model implementation, is essential.⁷⁹ FMs may also pose cybersecurity threats, and these tools could be exploited by malicious actors to extract sensitive patient data through techniques such as

Risks & Challenges for Implementing Foundation Models in Radiology

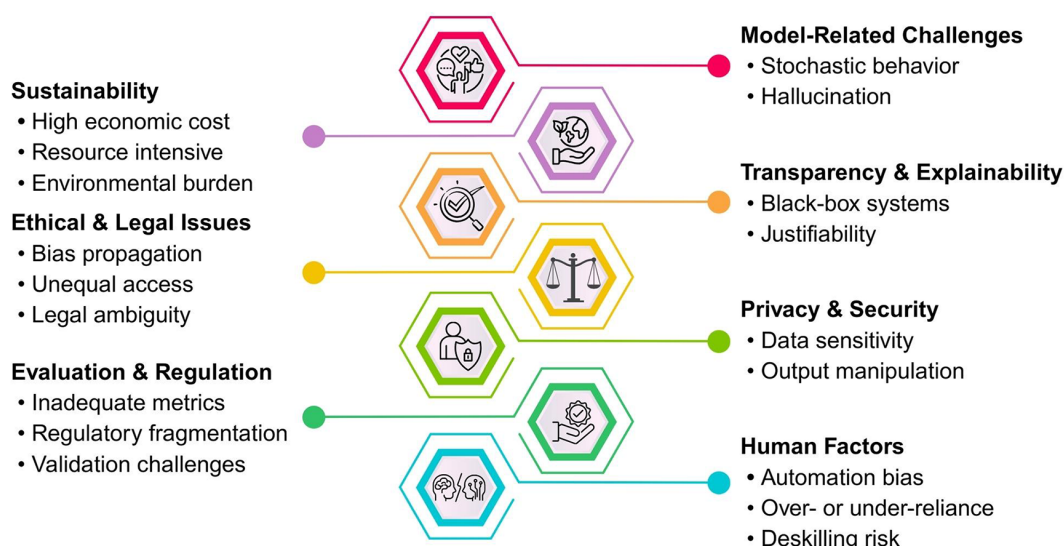


Figure 11. Potential risks and challenges ahead for implementing foundation models in radiology.

jailbreaking or to manipulate model output through techniques such as backdoor attacks.^{82,83} Ensuring robust security protocols and continuous monitoring is essential to safeguard patient data and maintain trust when implementing these tools in radiology.

Evaluating the performance of FMs presents another set of challenges. Traditional metrics, such as accuracy or F1 score, may be inadequate for assessing generative outputs or for evaluating the model's generation quality due to the lack of a reference standard.⁴⁷ Moreover, regulatory guidelines for ensuring the clinical safety and efficacy of these models are still in their infancy, and each country or region creating its own set of frameworks makes it harder to disseminate these tools (e.g., the AI Act across Europe and the Food and Drug Administration medical device law across the USA).⁸⁴ Rigorous validation processes and international regulatory alignment are necessary to overcome this hurdle. Besides the aforementioned challenges, foundational models introduce additional risks. For example, over-reliance on AI tools may lead to the deskilling of radiologists, weakening their ability to assess critically the AI-aided recommendations.⁸⁵ Ongoing education and training for radiologists are essential to mitigate deskilling and ensure appropriate use of these technologies.


Prospects

As previously discussed, attention-based FMs have represented a technological leap in

Environmental Impact



Energy consumption for generating a single AI image is roughly equivalent to half a smartphone's full battery charge.



Generating 10–50 medium-length chatbot responses may consume up to 0.5 liters of freshwater.

Figure 12. Environmental impact of foundation models. AI, artificial intelligence.

AI capabilities, with a wide range of potential applications in medical imaging. However, it should be noted that AI research is in continuous development, and even as LLMs

and FMs are just starting to be employed in the radiology domain, novel technologies are already aiming to complement or substitute current architectures and improve upon

their performance, alignment, and other limitations.

Among DL developments, state space models and recurrent neural networks currently represent promising architectures in the context of FMs, although the latter is a relatively mature technology.⁸⁶⁻⁸⁸ Both methods, with different specific implementations, including hybrid approaches,⁸⁹ incorporate recursive computation and representation of longer data sequences compared with “pure” transformer-based models. Furthermore, as is often the case in AI, these represent, at least in part, implementations of concepts initially theorized decades ago,⁹⁰ which have found new applications due to an increase in computational power and data availability.

Independent of the chosen neural network architecture, in the future, radiologists should expect (and increase their demand for) greater use of open-source software in the setting of generative AI and FMs. At the moment, this domain is largely dominated by proprietary (i.e., closed) technologies, which obfuscate the data used to train these tools, the network’s architecture, and the specific weights stored within the trained model. As previously mentioned, this lack of transparency represents a limitation to the implementation of FMs in healthcare, as well as running contrary to the principles contained within the European Union’s AI regulatory framework.⁹¹ Nevertheless, high-performance and large-scale FMs, which are also open-source, are already available, with Meta’s LLaMA being the most well-known. On the other hand, the increase in transparency afforded by open-source software comes with different considerations and potential tradeoffs. Although open-sourcing the model itself does not inherently compromise the privacy of the original training data (which are usually kept separate) or the input data, ensuring model security, preventing model

misuse, and mitigating potential risks such as cyberattacks requires careful governance; this may represent a significant issue in sensitive contexts, such as medical imaging. Furthermore, although open source does not represent an outright impediment to patenting, it does present a greater degree of challenge in protecting the technology behind a medical device and allowing a company to extract the economic value necessary to justify the large-scale investments required to develop such devices and the models running in the backend. This tension between private companies and public interest is not new to healthcare and has been the object of long debates in, for example, the setting of pharmaceuticals.⁹²⁻⁹⁴

AI and FMs can certainly look to these lessons to further establish the appropriate ethical and regulatory framework as these technologies increase their footprint in medical imaging, rather than attempting to reinvent the wheel. A clear sign of the relevance of these considerations is represented by the EU Commission’s recently announced intention to withdraw the proposed AI Liability Directive in its 2025 work program, demonstrating the regulator’s difficulties in balancing patient protection and incentivizing innovation.^{95,96}

AI and FMs will almost certainly impact healthcare, especially medical imaging,⁹⁷ in the future. In this setting, radiologists will need to be ready to increase their involvement in multidisciplinary teams. Deployment (and development) of FMs will require the expansion of the expertise requirements in imaging departments and closer collaboration with information technology, data science, and machine learning operations professionals. It could also be argued that the current vision in this profession regarding the implementation of this type of AI is still limited and mostly based on “adding on” FM to the current clinical workflow.⁹⁸ However,

it is also possible that this may not be the best strategy to implement this technology and may lead to unmet expectations and low impact on patient outcomes.^{99,100} Rather, the time may soon come to face the reality that FMs will require a radical rethinking of parts of the medical imaging practice: for example, regarding the scale of service delivery and role of the radiologist.¹⁰¹

Final thoughts

FMs represent a potential paradigm shift in AI, offering broad adaptability, multimodal integration, and improved generalizability across a wide range of tasks. In radiology, FMs have an immense potential to enable applications spanning image analysis, report generation, and integrative diagnostics across heterogeneous data sources. However, realizing this potential requires addressing key challenges, including issues of transparency, sustainability, data privacy, regulatory complexity, and ethical implementation. The inherent stochasticity and risk of bias in these models necessitate rigorous validation and continuous monitoring. Successful integration will require not only technical advancement but also adaptive clinical workflows and absolute transparency, potentially facilitated through open-source frameworks. Radiologists (along with other stakeholders) must play a central role in guiding the responsible development and deployment of FMs to ensure they augment, rather than undermine, the quality, safety, and equity of patient care. To this end, the authors of this international collaborative effort provide the radiology community with a set of practical recommendations based on the content extensively discussed in this work, to facilitate the better integration of FMs into clinical practice (Table 2). The authors hope that both the review and the accompanying recommendations will serve as a solid foundation for radiologists in adapting to rapidly evolving AI technologies, specifically FMs.

Table 2. Practical recommendations for radiologists and other healthcare stakeholders regarding the integration of artificial intelligence foundation models into clinical practice

Category	Recommendation
Safety and localization	<ul style="list-style-type: none">• Rigorously test foundation models on your data/protocols before any clinical use• Ensure strict patient privacy, anonymization, and regulatory compliance are provided by the vendor
Clinical use and trust	<ul style="list-style-type: none">• Prioritize transparent models (e.g., with explainable AI methods) to understand AI reasoning• Integrate AI thoughtfully to genuinely improve, not just add to, processes• Always retain final clinical judgment; AI assists, radiologists decide• Evaluate AI outputs in real-world settings (e.g., different scanners, different patient demographics) to assess robustness• Develop clear communication strategies for disclosing AI-assisted interpretations to referring physicians and patients
Responsible AI	<ul style="list-style-type: none">• Be watchful for potential biases and regularly assess following updates for fairness• Demand vendor transparency• Push for third-party validation of AI tools (not just vendor-reported metrics)• Participate in and support open science initiatives for shared model evaluation• Be mindful of AI's environmental/computational, economic, and social footprint
Future-proofing	<ul style="list-style-type: none">• Continuously learn AI principles, data science, and ethics to build AI literacy to help make informed decisions• Follow evolving AI regulations (e.g., AI Act) and best practices• Build internal AI governance structures or committees to oversee and audit ongoing usage and updates

AI, artificial intelligence.

Acknowledgements

ChatGPT (4o) and Gemini 2.5 Pro were partially used under strict author supervision for language refinement, summarization, and brainstorming to enhance the content originally written by the authors. CB received research support, independent of this work, from Promedica Foundation, Chur, CH.

Footnotes

Conflict of interest disclosure

Tuğba Akıncı D'Antonoli and Burak Koçak are Section Editors in Diagnostic and Interventional Radiology (Section: Artificial Intelligence and Informatics). They had no involvement in the peer review of this article and had no access to information regarding its peer review.

References

1. Litjens G, Kooi T, Bejnordi BE, et al. A survey on deep learning in medical image analysis. *Med Image Anal.* 2017;42:60-88. [\[Crossref\]](#)
2. Aggarwal R, Sounderajah V, Martin G, et al. Diagnostic accuracy of deep learning in medical imaging: a systematic review and meta-analysis. *NPJ Digital Med.* 2021;4(1):65. [\[Crossref\]](#)
3. Bommasani R, Hudson DA, Adeli E, et al. On the opportunities and risks of foundation models. *arXiv.* 2021. [\[Crossref\]](#)
4. Azad B, Azad R, Eskandari S, et al. Foundational models in medical imaging: a comprehensive survey and future vision. *arXiv.* 2023. [\[Crossref\]](#)
5. Zhou C, Li Q, Li C, et al. A comprehensive survey on pretrained foundation models: a history from BERT to ChatGPT. *arXiv*, v2. 1 May 2023. [\[Crossref\]](#)
6. Myers D, Mohawesh R, Chellaboina VI, et al. Foundation and large language models: fundamentals, challenges, opportunities, and social impacts. *Cluster Comput.* 2024;27(1):1-26. [\[Crossref\]](#)
7. Akinci D'Antonoli T, Stanzione A, Bluethgen C, et al. Large language models in radiology: fundamentals, applications, ethical considerations, risks, and future directions. *Diagn Interv Radiol.* 2024;30(2):80-90. [\[Crossref\]](#)
8. Li C, Gan Z, Yang Z, et al. Multimodal foundation models: from specialists to general-purpose assistants. *arXiv.* 2023. [\[Crossref\]](#)
9. Fei N, Lu Z, Gao Y, et al. Towards artificial general intelligence via a multimodal foundation model. *Nat Commun.* 2022;13(1):3094. [\[Crossref\]](#)
10. Simon BD, Ozyoruk KB, Gelikman DG, Harmon SA, Türkbey B. The future of multimodal artificial intelligence models for integrating imaging and clinical metadata: a narrative review. *Diagn Interv Radiol.* October 1, 2024. [\[Crossref\]](#)
11. Acosta JN, Falcone GJ, Rajpurkar P, Topol EJ. Multimodal biomedical AI. *Nat Med.* 2022;28(9):1773-1784. [\[Crossref\]](#)
12. Tu T, Azizi S, Driess D, et al. Towards generalist biomedical AI. *NEJM AI.* 2024;1(3). [\[Crossref\]](#)
13. Zhang S, Xu Y, Usuyama N, et al. A multimodal biomedical foundation model trained from fifteen million image-text pairs. *NEJM AI.* 2025;2(1). [\[Crossref\]](#)
14. Strotzer QD, Nieberle F, Kupke LS, et al. Toward foundation models in radiology? Quantitative assessment of GPT-4V's multimodal and multianatomic region capabilities. *Radiology.* 2024;313(2):240955. [\[Crossref\]](#)
15. Busch F, Hoffmann L, Dos Santos DP, et al. Large language models for structured reporting in radiology: past, present, and future. *Eur Radiol.* 2025;35(5):2589-2602. [\[Crossref\]](#)
16. Meddeb A, Lükens S, Busch F, et al. Large language model ability to translate CT and MRI free-text radiology reports into multiple languages. *Radiology.* 2024;313(3):241736. [\[Crossref\]](#)
17. Fink MA, Bischoff A, Fink CA, et al. Potential of ChatGPT and GPT-4 for data mining of free-text CT reports on lung cancer. *Radiology.* 2023;308(3):231362. [\[Crossref\]](#)
18. Patil NS, Huang RS, van der Pol CB, Larocque N. Comparative performance of ChatGPT and bard in a text-based radiology knowledge assessment. *Can Assoc Radiol J.* 2024;75(2):344-350. [\[Crossref\]](#)
19. Chen M, Zhang M, Yin L, et al. Medical image foundation models in assisting diagnosis of brain tumors: a pilot study. *Eur Radiol.* 2024;34(10):6667-6679. [\[Crossref\]](#)
20. Ali L, Alnajjar F, Swavaf M, Elharrouss O, Abd-Alrazaq A, Damseh R. Evaluating segment anything model (SAM) on MRI scans of brain tumors. *Sci Rep.* 2024;14(1):21659. [\[Crossref\]](#)
21. Wu C, Zhang X, Zhang Y, Wang Y, Xie W. Towards generalist foundation model for radiology. *arXiv.* 2023. [\[Crossref\]](#)
22. Krishnan R, Rajpurkar P, Topol EJ. Self-supervised learning in medicine and healthcare. *Nat Biomed Eng.* 2022;6(12):1346-1352. [\[Crossref\]](#)
23. Huang S-C, Pareek A, Jensen M, Lungren MP, Yeung S, Chaudhari AS. Self-supervised learning for medical image classification: a systematic review and implementation guidelines. *npj Digital Med.* 2023;6(1):74. [\[Crossref\]](#)
24. Hoffmann J, Borgeaud S, Mensch A, et al. Training compute-optimal large language models. *arXiv.* 2022. [\[Crossref\]](#)
25. Ouyang L, Wu J, Jiang X, et al. Training language models to follow instructions with human feedback. *arXiv.* 2022. [\[Crossref\]](#)
26. Liu H, Li C, Wu Q, Lee YJ. Visual Instruction Tuning. *Advances in Neural Information*

- Processing Systems*. December 15, 2023. [\[Crossref\]](#)
27. Schaeffer R, Miranda B, Koyejo S. Are emergent abilities of large language models a mirage? *arXiv*. 2023. [\[Crossref\]](#)
 28. Thrun S, Pratt L. Learning to learn: introduction and overview. In: Thrun S, Pratt L, eds. *Learning to Learn*. Boston, MA: Springer;1998. [\[Crossref\]](#)
 29. Caruana R. Multitask Learning. *Springer Science and Business Media LLC*. 1997;28:41-75. [\[Crossref\]](#)
 30. Raina R, Ng AY, Koller D. Constructing informative priors using transfer learning. In: *Proceedings of the 23rd International Conference on Machine Learning - ICML '06*. ACM Press; 2006:713-720. [\[Crossref\]](#)
 31. Wu P, Dietterich TG. Improving SVM accuracy by training on auxiliary data sources. In: *Twenty-First International Conference on Machine Learning - ICML '04*. ACM Press; 2004:110. [\[Crossref\]](#)
 32. Palatucci M, Pomerleau D, Hinton GE, Mitchell TM. Zero-shot learning with semantic output codes. *Adv Neural Inf Process Syst*. 2009. [\[Crossref\]](#)
 33. Kojima T, Gu SS, Reid M, Matsuo Y, Iwasawa Y. Large language models are zero-shot reasoners. *arXiv*. 2022. [\[Crossref\]](#)
 34. Brown TB, Mann B, Ryder N, et al. Language models are few-shot learners. *arXiv*. 2020. [\[Crossref\]](#)
 35. Oquab M, Darcet T, Moutakanni T, et al. DINOv2: learning robust visual features without supervision. *arXiv*. 2023. [\[Crossref\]](#)
 36. Pai S, Hadzic I, Bontempi D, et al. Vision foundation models for computed tomography. *arXiv*. 2025. [\[Crossref\]](#)
 37. Tiu E, Talus E, Patel P, Langlotz CP, Ng AY, Rajpurkar P. Expert-level detection of pathologies from unannotated chest X-ray images via self-supervised learning. *Nat Biomed Eng*. 2022;6(12):1399-1406. [\[Crossref\]](#)
 38. Chen Z, Varma M, Delbrouck JB, et al. CheXagent: Towards a foundation model for chest X-ray interpretation. *arXiv*. 2024. [\[Crossref\]](#)
 39. Deperrois N, Matsuo H, Ruipérez-Campillo S, et al. RadVLM: a multitask conversational vision-language model for radiology. *arXiv*. 2025. [\[Crossref\]](#)
 40. Bluethgen C, Chambon P, Delbrouck JB, et al. A vision-language foundation model for the generation of realistic chest X-ray images. *Nat Biomed Eng*. 2025;9(4):494-506. [\[Crossref\]](#)
 41. Blankemeier L, Cohen JP, Kumar A, et al. Merlin: a vision language foundation model for 3D computed tomography. *arXiv*. 2024. [\[Crossref\]](#)
 42. Hamamci IE, Er S, Wang C, et al. Developing generalist foundation models from a multimodal dataset for 3D computed tomography. *arXiv*. 2024. [\[Crossref\]](#)
 43. Sun Y, Wang L, Li G, Lin W, Wang L. A foundation model for enhancing magnetic resonance images and downstream segmentation, registration and diagnostic tasks. *Nat Biomed Eng*. 2025;9(4):521-538. [\[Crossref\]](#)
 44. Vaswani A, Shazeer N, Parmar N, et al. Attention is all you need. *arXiv*. 2017. [\[Crossref\]](#)
 45. Radford A, Kim JW, Hallacy C, et al. Learning transferable visual models from natural language supervision. *arXiv*. 2021. [\[Crossref\]](#)
 46. Dosovitskiy A, Beyer L, Kolesnikov A, et al. An image is worth 16x16 words: transformers for image recognition at scale. *arXiv*. 2020. [\[Crossref\]](#)
 47. Paschali M, Chen Z, Blankemeier L, et al. Foundation models in radiology: what, how, why, and why not. *Radiology*. 2025;314(2):e240597. [\[Crossref\]](#)
 48. Han Z, Gao C, Liu J, Zhang J, Zhang SQ. Parameter-efficient fine-tuning for large models: a comprehensive survey. *arXiv*. 2024. [\[Crossref\]](#)
 49. Zhang S, Metaxas D. On the challenges and perspectives of foundation models for medical image analysis. *Med Image Anal*. 2024;91:102996. [\[Crossref\]](#)
 50. Rayed E, Islam SMS, Niha SI, Jim JR, Kabir M, Mridha MF. Deep learning for medical image segmentation: State-of-the-art advancements and challenges. *Inform Med Unlocked*. 2024;47:101504. [\[Crossref\]](#)
 51. Akinci D'Antonoli T, Berger LK, Indrakanti AK, et al. TotalSegmentator MRI: Robust Sequence-independent Segmentation of Multiple Anatomic Structures in MRI. *Radiology*. 2025;314(2):e241613. [\[Crossref\]](#)
 52. Wasserthal J, Breit HC, Meyer MT, et al. Totalsegmentator: robust segmentation of 104 anatomic structures in CT images. *Radiol Artif Intell*. 2023;5(5):e230024. [\[Crossref\]](#)
 53. Kirillov A, Mintun E, Ravi N, et al. Segment anything. *arXiv*. 2023. [\[Crossref\]](#)
 54. He S, Bao R, Li J, et al. Computer-Vision Benchmark Segment-Anything Model (SAM) in Medical Images: Accuracy in 12 Datasets. *arXiv*. 2023. [\[Crossref\]](#)
 55. Shi P, Qiu J, Abaxi SMD, Wei H, Lo FP-W, Yuan W. Generalist vision foundation models for medical imaging: a case study of Segment Anything Model on zero-shot medical segmentation. *Diagnostics (Basel)*. 2023;13(11). [\[Crossref\]](#)
 56. Chen S, Ma K, Zheng Y. Med3D: Transfer learning for 3D medical image analysis. *arXiv*. 2019. [\[Crossref\]](#)
 57. Zhu J, Hamdi A, Qi Y, Jin Y, Wu J. Medical SAM 2: segment medical images as video via Segment Anything Model 2. *arXiv*. 2024. [\[Crossref\]](#)
 58. Butoi VI, Ortiz JGG, Ma T, Sabuncu MR, Guttag J, Dalca AV. UniverSeg: Universal medical image segmentation. *arXiv*. 2023. [\[Crossref\]](#)
 59. Xu L, Ni Z, Sun H, Li H, Zhang S. A foundation model for generalizable disease diagnosis in chest X-ray images. *arXiv*. 2024. [\[Crossref\]](#)
 60. Yan A, McAuley J, Lu X, et al. RadBERT: adapting transformer-based language models to radiology. *Radiol Artif Intell*. 2022;4(4):e210258. [\[Crossref\]](#)
 61. Chen Z, Song Y, Chang TH, Wan X. Generating radiology reports via memory-driven transformer. *arXiv*. 2020. [\[Crossref\]](#)
 62. Luo L, Vairavamurthy J, Zhang X, et al. ReXplain: translating radiology into patient-friendly video reports. *arXiv*. 2024. [\[Crossref\]](#)
 63. Chen Q, Hu X, Wang Z, Hong Y. MedBLIP: Bootstrapping language-image pre-training from 3D medical images and texts. *arXiv*. 2023. [\[Crossref\]](#)
 64. Moor M, Banerjee O, Abad ZSH, et al. Foundation models for generalist medical artificial intelligence. *Nature*. 2023;616(7956):259-265. [\[Crossref\]](#)
 65. He Y, Huang F, Jiang X, et al. Foundation model for advancing healthcare: challenges, opportunities, and future directions. *arXiv*. 2024. [\[Crossref\]](#)
 66. Wolf D, Payer T, Lissin CS, et al. Self-supervised pre-training with contrastive and masked autoencoder methods for dealing with small datasets in deep learning for medical imaging. *Sci Rep*. 2023;13(1):20260. [\[Crossref\]](#)
 67. Koçak B, Ponsiglione A, Stanzione A, et al. Bias in artificial intelligence for medical imaging: fundamentals, detection, avoidance, mitigation, challenges, ethics, and prospects. *Diagn Interv Radiol*. 2025;31(2):75-88. [\[Crossref\]](#)
 68. Tripathi S, Gabriel K, Dheer S, et al. Understanding biases and disparities in radiology AI datasets: a review. *J Am Coll Radiol*. 2023;20(9):836-841. [\[Crossref\]](#)
 69. Sizikova E, Badal A, Delfino JG, et al. Synthetic data in radiological imaging: current state and future outlook. *BJR Artificial Intelligence*. 2024;1(1). [\[Crossref\]](#)
 70. Chambon P, Bluethgen C, Delbrouck JB, et al. RoentGen: Vision-language foundation model for chest X-ray generation. *arXiv*. 2022. [\[Crossref\]](#)
 71. He Y, Huang F, Jiang X, et al. Foundation model for advancing healthcare: challenges, opportunities and future directions. *IEEE Rev Biomed Eng*. 2025;18:172-191. [\[Crossref\]](#)
 72. Cai P, Yao Z, Liu F, et al. PaniniQA: enhancing patient education through interactive question answering. *arXiv*. 2023. [\[Crossref\]](#)
 73. Kung TH, Cheatham M, Medenilla A, et al. Performance of ChatGPT on USMLE: potential for AI-assisted medical education using large language models. *PLOS Digit Health*. 2023;2(2):e0000198. [\[Crossref\]](#)

74. Sun K, Xue S, Sun F, et al. Medical multimodal foundation models in clinical diagnosis and treatment: applications, challenges, and future directions. *arXiv*. 2024. [\[Crossref\]](#)
75. Luccioni S, Jernite Y, Strubell E. Power hungry processing: watts driving the cost of AI deployment? In: *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. ACM; 2024:85-99. [\[Crossref\]](#)
76. Li P, Yang J, Islam MA, Ren S. Making AI Less "Thirsty": uncovering and addressing the secret water footprint of AI models. *arXiv*. 2023. [\[Crossref\]](#)
77. Doo FX, Vosschenrich J, Cook TS, et al. Environmental sustainability and ai in radiology: a double-edged sword. *Radiology*. 2024;310(2):e232030. [\[Crossref\]](#)
78. Kocak B, Ponsiglione A, Romeo V, Ugga L, Huisman M, Cuocolo R. Radiology AI and sustainability paradox: environmental, economic, and social dimensions. *Insights Imaging*. 2025;16(1):88. [\[Crossref\]](#)
79. Akinci D'Antonoli T. Ethical considerations for artificial intelligence: an overview of the current radiology landscape. *Diagn Interv Radiol*. 2020;26(5):504-511. [\[Crossref\]](#)
80. Bhattacharjya D, Lee J, Agravante DJ, Ganesan B, Marinescu R. Foundation model sherpas: guiding foundation models through knowledge and reasoning. *arXiv*. 2024. [\[Crossref\]](#)
81. Gaube S, Suresh H, Raue M, et al. Do as AI say: susceptibility in deployment of clinical decision-aids. *NPJ Digit Med*. 2021;4(1):31. [\[Crossref\]](#)
82. Yao Y, Duan J, Xu K, Cai Y, Sun Z, Zhang Y. A survey on large language model (LLM) security and privacy: the good, the bad, and the ugly. *High-Confidence Computing*. 2024;100211. [\[Crossref\]](#)
83. Ponsiglione A, Cuocolo R. Radiology under siege? Adversarial attacks against deep learning algorithms. *Eur J Radiol*. 2023;169:111156. [\[Crossref\]](#)
84. Aboy M, Minssen T, Vayena E. Navigating the EU AI Act: implications for regulated digital medical products. *NPJ Digital Med*. 2024;7(1):237. [\[Crossref\]](#)
85. Dratsch T, Chen X, Rezazade Mehrizi M, et al. Automation bias in mammography: the impact of artificial intelligence BIRADS suggestions on reader performance. *Radiology*. 2023;307(4):e222176. [\[Crossref\]](#)
86. Sieber J, Alonso CA, Didier A, Zeilinger MN, Orvieto A. Understanding the differences in foundation models: attention, state space models, and recurrent neural networks. *arXiv*. 2024. [\[Crossref\]](#)
87. Gu A, Goel K, Ré C. Efficiently modeling long sequences with structured state spaces. *arXiv*. 2021. [\[Crossref\]](#)
88. Stanić A, Ashley D, Serikov O, et al. The languini kitchen: enabling language modelling research at different scales of compute. *arXiv*. 2023. [\[Crossref\]](#)
89. De S, Smith SL, Fernando A, et al. Griffin: mixing gated linear recurrences with local attention for efficient language models. *arXiv*. 2024. [\[Crossref\]](#)
90. Hochreiter S, Schmidhuber J. Long short-term memory. *Neural Comput*. 1997;9(8):1735-1780. [\[Crossref\]](#)
91. Kotter E, D'Antonoli TA, Cuocolo R, et al. Guiding AI in radiology: ESR's recommendations for effective implementation of the European AI Act. *Insights Imaging*. 2025;16(1):33. [\[Crossref\]](#)
92. Statman M, Tyebjee TT. Trademarks, patents, and innovation in the ethical drug industry. *J Mark*. 1981;45(3):71-81. [\[Crossref\]](#)
93. De George RT. Intellectual property and pharmaceutical drugs: an ethical analysis. *Bus Ethics Q*. 2005;15(4):549-575. [\[Crossref\]](#)
94. Millum J. Are pharmaceutical patents protected by human rights? *J Med Ethics*. 2008;34(11):e25. [\[Crossref\]](#)
95. Ruscemeier H. The de-regulatory turn of the EU Commission. *Verfassungsblog*. 2025. [\[Crossref\]](#)
96. VOSS A. Carriages preview | Legislative Train Schedule. Accessed April 26, 2025. [\[Crossref\]](#)
97. Kocak B, Baessler B, Cuocolo R, Mercaldo N, Pinto Dos Santos D. Trends and statistics of artificial intelligence and radiomics research in radiology, nuclear medicine, and medical imaging: bibliometric analysis. *Eur Radiol*. 2023;33(11):7542-7555. [\[Crossref\]](#)
98. Linguraru MG, Bakas S, Aboian M, et al. Clinical, cultural, computational, and regulatory considerations to deploy AI in radiology: perspectives of RSNA and MICCAI experts. *Radiol Artif Intell*. 2024;6(4):e240225. [\[Crossref\]](#)
99. Del Gaizo AJ, Osborne TF, Shahoumian T, Sherrier R. Deep learning to detect intracranial hemorrhage in a national teleradiology program and the impact on interpretation time. *Radiol Artif Intell*. 2024;6(5):e240067. [\[Crossref\]](#)
100. Savage CH, Tanwar M, Elkassem AA, et al. Prospective evaluation of artificial intelligence triage of intracranial hemorrhage on noncontrast head CT examinations. *AJR Am J Roentgenol*. 2024;223(5):e2431639. [\[Crossref\]](#)
101. Pierre K, Haneberg AG, Kwak S, et al. Applications of artificial intelligence in the radiology roundtrip: process streamlining, workflow optimization, and beyond. *Semin Roentgenol*. 2023;58(2):158-169. [\[Crossref\]](#)