JOURNAL OF CULTURAL ANALYSIS AND SOCIAL CHANGE

Journal of Cultural Analysis and Social Change,
2026, 11(2), 193-207
ISSN: 2589-1316

Interdisciplinary Synergy Between Radiologist and ML Engineer for augmented Whole-Body MRI Interpretation in the AI Era: From Pixels to Decisions

Christos B. Zachariadis[1,2*], Charalampos Z. Patrikakis[3], George Georgoudis[4], Helen C. Leligou[1]

[1] *Department of Industrial Design and Production Engineering, University of West Attica, P. Ralli & Thivon 250, 12244 Athens, Greece*
[2] *First Department of Orthopaedic Surgery, School of Medicine, National and Kapodistrian University of Athens, "Attikon" University General Hospital, Athens, Greece*
[3] *Department of Electrical and Electronics Engineering, University of West Attica, P. Ralli & Thivon 250, 12244 Athens, Greece*
[4] *Department of Physiotherapy, University of West Attica, P. Ralli & Thivon 250, 12244 Athens, Greece*

*__Corresponding Author:__ drzachariadis@uniwa.gr

Citation: Zachariadis, C. B., Patrikakis, C. Z, Georgoudis, H and Leligou, H. C. (2026). Interdisciplinary Synergy Between Radiologist and ML Engineer for augmented Whole-Body MRI Interpretation in the AI Era: From Pixels to Decisions, *Journal of Cultural Analysis and Social Change*, 11(2), 193-207. https://doi.org/10.64753/jcasc.v11i2.4872

Published: May 25, 2026

ABSTRACT

The concept of an AI-assisted comprehensive, accessible, and safe diagnostic modality has gained significant attention in recent years, reflecting the evolving needs and expectations of contemporary society. Current advancements in medical imaging have enabled high-resolution, detailed, non-invasive whole body image reconstruction, allowing for more accurate lesion detection and characterization and, therefore, facilitating more effective, targeted treatment. Practical employment of such solutions can become feasible and clinically useful through several key technologies that leverage state of the art adaptive algorithms. From this perspective, we describe the theoretical principles of a conceptual diagnostic network for AI-supported whole-body (WB) MRI analysis. We propose a framework that integrates well-established and widely studied AI tools, including Convolutional Neural Networks for image analysis, Natural Language Generation for structured reporting and Federated Learning for decentralized model training, into a multimodal diagnostic tool. These elements are selected for their strong academic maturity, broad validation across medical imaging studies, and proven compatibility with modern WB-MRI scanners. The expert-in-the-loop precept places the radiologist at the helm of this system, where they maintain clinical oversight and accountability for all decision-making, while also contributing to the ongoing readjustment of the functional algorithm. Inherently dependent on actual data availability, the complexity of the process raises issues related to patient psychological burden, as well as legality, fairness, and ethical governance.

Keywords: AI-assisted diagnosis; Artificial Intelligence; semi-automated; whole-body MRI; convolutional neural network; natural language generation; federated learning; explainable AI; generative AI; anonymization; encryption.

INTRODUCTION

A comprehensive whole-body radiological examination capable of routinely detecting a broad spectrum of pathological entities represents an ambitious aspiration of present-day society. Nonetheless, ongoing developments in medical imaging and state-of-the-art adaptive algorithms are gradually transforming the notion from science fiction to applicable reality [1-3].

The primary modality that renders this concept viable is the whole-body magnetic resonance imaging (WB-MRI), when integrating technologies that add functional and compositional information beyond standard anatomical analysis. Unlike traditional scanning tools that may either focus on limited regional representation, or lack the capacity for detailed depiction, WB-MRI can facilitate head-to-heel, ionizing radiation-free, high-resolution 3D image reconstruction in a single session [4-6]. While its use in wider population screening is not yet indicated under current protocols, certain diagnostic guidelines include it as a recommendation for their respective conditions; it is increasingly applied in neoadjuvant staging, postoperative surveillance, and evaluation of treatment response in oncological patients [7,8].

Two major constraints impede the process: the massive volume of acquisition data to be processed, and the time-intensive task of clinical reporting. In the context of WB-MRI implementation, the emerging application of convolutional neural networks (CNNs) has enabled automated image segmentation, reconstruction and denoising, as well as lesion detection, turning raw input MRI signals into detailed anatomical and pathological maps. Through natural language generation (NLG) processes, the spatial coordinates and quantitative data are converted into a structured diagnostic draft. Despite their limitations, such multimodal constructs hold immense potential of becoming a valuable assistive force in daily medical practice. Indeed, several studies have demonstrated the compatibility of MRI data with classification procedures performed by adaptive algorithms [5,6].

The present proposition is composed following the premise that interaction between the radiologist and an AI-enhanced image and language processing component achieves improved diagnostic accuracy, speed in delivery of the official report, and expert's satisfaction. The article outlines the key operational requirements for a seamless workflow within a geographically distributed AI-assisted diagnostic system that uses WB-MRI; it summarizes the underlying currently available technologies that constitute its core [2,9], and presents the detailed interplay between the machine-learning system and the medical professional, who has a clearly defined role in evaluating regions of interest, resolving local inconsistencies, and producing the final report within the pipeline. It further discusses possible psychological implications on the examinee, and addresses issues pertaining to data privacy, legal compliance, algorithmic fairness, and ethical responsibility.

Although this is not a review article, each parameter was informed by a literature search using non-systematic queries in major scientific databases. The overall approach does not aim to delve into the technical details of advanced AI technologies but rather to delineate human–machine cooperation within a harmonized framework of distinct and complementary roles.

The hypothesis

While advances in AI-enhanced architectures for medical imaging interpretation have been widely utilized in segmentation and classification tasks [1,5,10,11], their coordinated integration within a unified geographically dispersed network has not been examined in comparative studies. Access to a broader spectrum of data enhances machine-learning accuracy by enabling the model to learn from a more wide-ranging, representative, and comprehensive set of examples [12]; hence a processing system that leverages federated learning (FL) technologies has the potential to ensure diversity and abundance of raw data, while maintaining efficiency in resource deployment. Federated Learning (FL) can enhance model generalization by enabling exposure to diverse datasets without compromising privacy [12].

In the proposed setting, the system aligns three elements—the WB-MRI modality, the adaptive algorithm, and the radiologist —forming a tripartite structure. The hypothesized workflow aims for (a) reduced turnaround time from scan to report, (b) higher diagnostic sensitivity through decreased false negatives, and (c) improved radiologist satisfaction, expressed as willingness to continue using the platform. An additional objective is the generation of high-quality, human validated training datasets, for iterative model refinement.

As an initial clinical use case, the framework could target whole-body MRI in multiple myeloma, utilizing standardized protocols, e.g. MY-RADS, for lesion detection and treatment-response evaluation. The AI output would include lesion segmentation maps and a structured draft report summarizing the skeletal disease burden. The radiologist in charge would review these outputs, correct discrepancies, and validate the results. Accuracy, speed, and user acceptance will serve as early validation endpoints. (Figure 1).

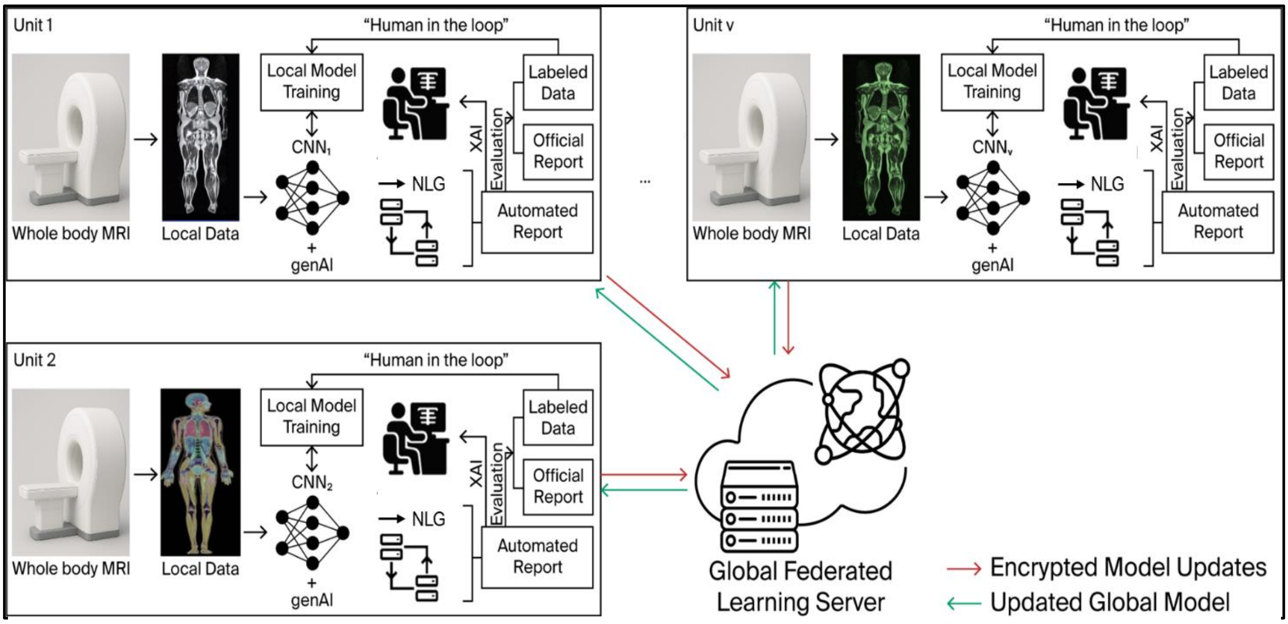

Figure 1. Architecture of the proposed network. Following WB-MRI examination, generated data is processed by a hybrid CNN - NLG model, predict a preliminary diagnosis, which the radiologist will evaluate, communicate to the patient, and upload for model training and periodical federated sharing.

To evaluate the innovation of our hypothesis, a narrative review was conducted. A literature search of all articles from January 2020 to April 2026 published in English was undertaken using Google Scholar and PubMed–Medline. Search criteria included the terms "Whole Body MRI" AND "Convolutional Neural Network" AND "Federated Learning" mentioned anywhere in the title, abstract or full text of the articles. The search resulted in 47 unique articles, which were examined for conceptual and methodological similarities to the proposed federated whole-body MRI network. While CNN and FL based approaches for automated diagnostic tasks in medical imaging are increasingly reported, no comparable integrated platform has been identified in the current literature

Network overview

The initial operation of the network will involve the central training of a CNN model. Among various CNN architectures employed for medical image interpretation, VGGNet, ResNet, and U-net currently exhibit convenient characteristics, and considerable popularity [13-15]. At this stage, available datasets from archives of the cooperating radiological units, combined with existing data repositories [16-18] will facilitate the inception of model training. The procedure should comply with data protection regulations. The occurring baseline model will be transmitted to the local subsets.

Up to a certain point, each participating radiological facility will constitute an autonomous processing unit. Within the pipeline, the local system will conduct a WB-MRI examination on patients whose management requires whole-body scanning either for staging their disease or for assessing response to therapy. From this perspective, the initial target population will focus on patients with conditions for which standardized WB-MRI protocols have clearly been developed. These conditions currently include multiple myeloma, for which the Myeloma Response Assessment and Diagnosis System (MY-RADS) was published in 2019, and advanced prostate cancer, for which the recommended WB-MRI specifications have recently been defined in the Metastasis Reporting and Data System for Prostate Cancer (MET-RADS-P) guidelines [8]. Alternatively, initial implementation could target screening of a specific high-risk population, i.e. individuals with cancer predisposition syndromes, following the Oncologically Relevant Findings Reporting and Data System (ONCO-RADS) guidelines [4,8]. A potential later expansion of indications is discussed in the WB-MRI section.

The WB-MRI scanner will yield highly detailed images of both bone and soft tissues —adipose, connective, muscular, neural, vascular —that will serve as original input data for the local CNN. The deep learning algorithm will scan the acquisitions immediately after they are uploaded, will classify them, flag areas whose physical properties deviate from normal patterns, and map them onto a 3D anatomical model as potential lesions. The findings will then be passed onto a natural language generation (NLG) module, that will read the provided evidence and convert it into a clinically readable, structured report. An explainable-AI (XAI) component will record and expose information, such as model decision rationale based on image quality metrics or sequence confidence, to enhance transparency and interpretability of the output.

The product reaching the radiologist will consist of the original input slices, a set of tagged slices demarcating areas of interest, and a preliminary textualized diagnostic assessment. This information will subsequently be reviewed, corrected or expanded, and validated by the expert. The involvement of the human element will ensure the reliability and accuracy of the conclusive diagnosis, before the clinical decision is made and the actual finalized report is delivered to the patient. The workflow from Wb-MRI examination to official report is depicted in Figure 2.

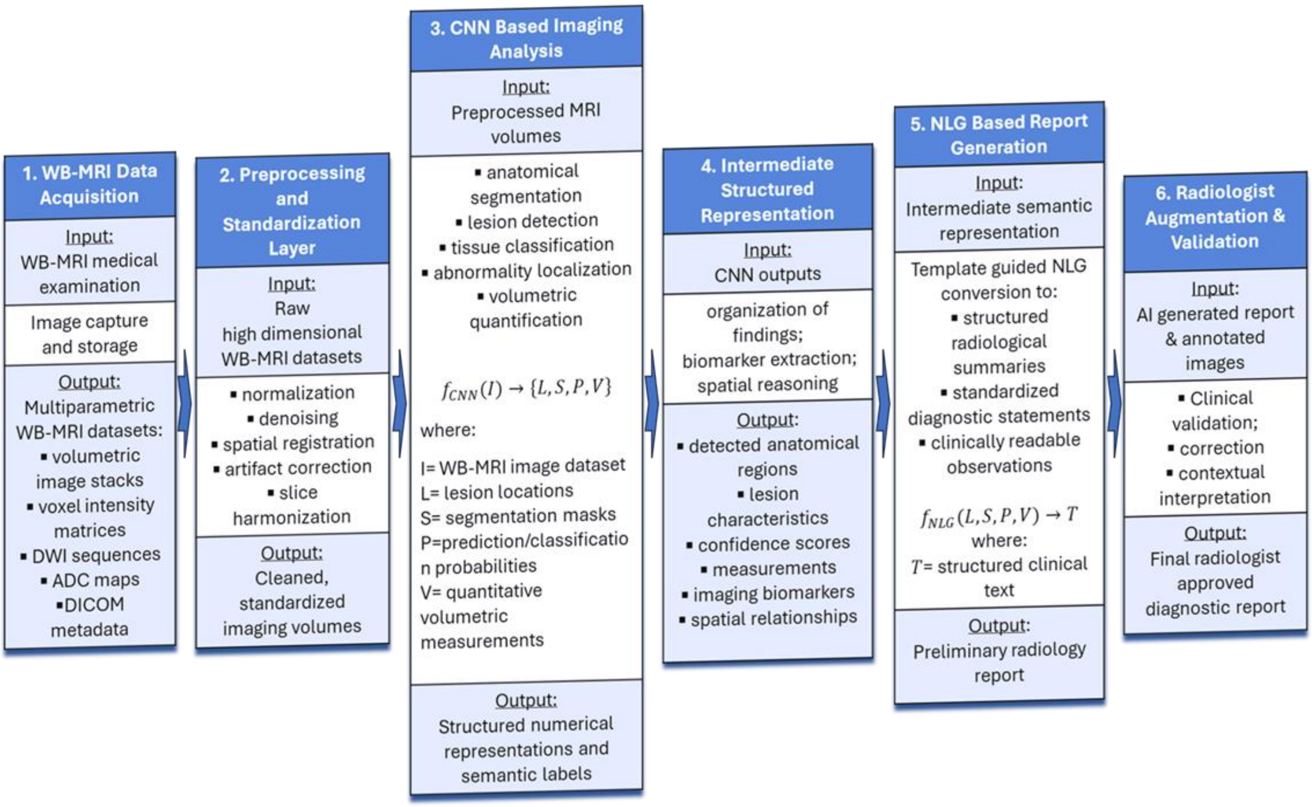

Figure 2. ***The workflow of the WB-MRI interpretation pipeline from raw data to official medical report.***

Following anonymization and encryption, the official radiology report will be converted into labels for supervised training of the local CNN and NLG models (Figure 3). The resulting datasets could also be stored and made available for training of other external systems (outside the network), under careful preset conditions. At predefined regular intervals, or after a specified number of training epochs, only model updates will be transmitted to the central FL server, rather than raw medical images or patient-level information. These updates will be securely aggregated, and used to produce a refined primary model capable of capturing and processing shared imaging features across all the cooperating facilities, while preserving confidentiality and maintaining compliance with regulatory requirements. After collecting updates from all collaborating units, the upgraded model will be redistributed back to the participating sites, where it will undergo successive cycles of local training, inference, and clinical decision-support. Through this iterative process, the system will contribute to gradually improving the accuracy of the finalized report document [12].

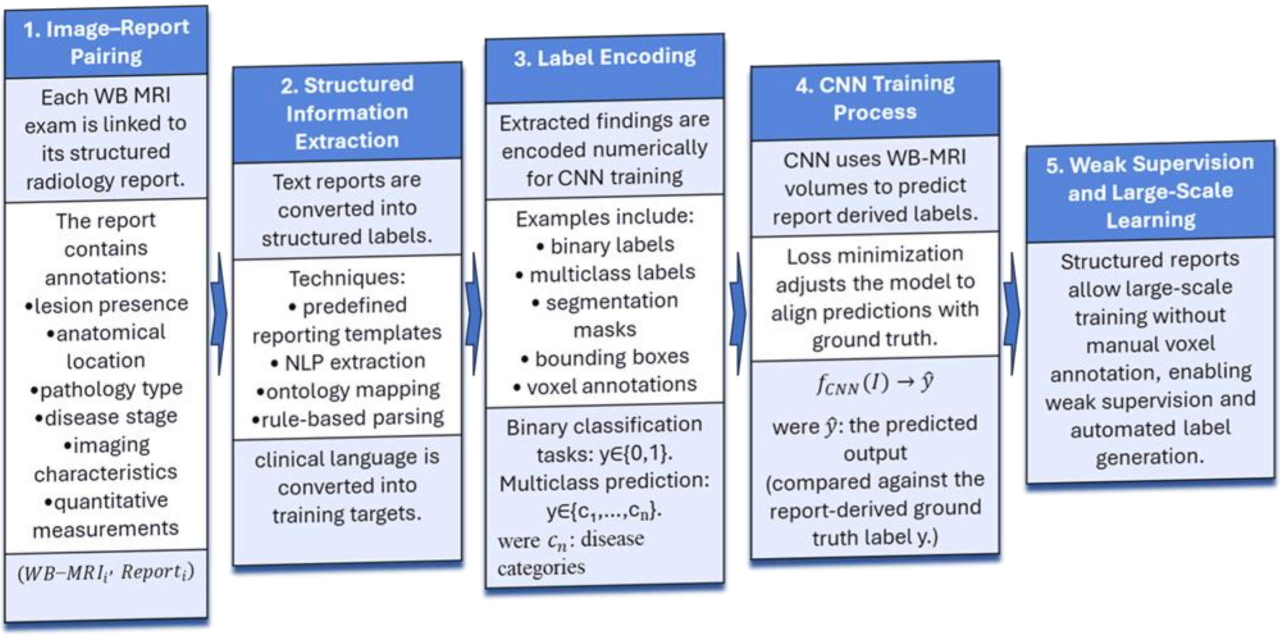

*Figure 3. **Structured Radiology Reports for CNN training.***

The nature of the design mandates that the contributing healthcare centers be equipped with state-of-the-art WB-MRI scanners, desirably open or semi-open to improve comfort of the examinee, reducing claustrophobia. The radiologist assigned with the evaluation, refinement and approval of the final account should be fully informed and consent to their role, while also being accustomed to interacting with the employed AI tools. In a real-world setting, and in order to minimize heterogeneity of raw data due to geographic disparity and differences in scanner specifications, a suitable candidate to test the first operational use case could be a single established private network of diagnostic laboratories that would agree to include pilot use of the proposed platform in its daily workload of MRI exams.

To reflect practical and realistic deployment of the platform, a three-stage implementation rationale can be adopted. In the first stage (Pilot Implementation), a single-center or small-network can be utilized within a controlled diagnostic laboratory setting, focusing on a specific diagnostic indication, e.g. multiple myeloma. At this stage, the core components will include the CNN and NLG modules, as well as basic XAI tools, aiming to assess workflow integration, diagnostic accuracy, and radiologist satisfaction. After successful proof-of-concept validation, guided by the performance outcomes of the first stage, a second stage (Scaling) will allow the system to expand across multiple healthcare facilities, incorporating federated learning to enable privacy-preserving model generalization and cross-center standardization through progressively larger datasets. In the third stage (Future Expansion), a research-focused phase will explore the integration of generative AI, preferably GAN-based approaches [19], and blockchain—as complementary, optional components; the goal, at that point, will be data augmentation, auditability, and long-term trust infrastructure. This staged stratification, emphasing practical feasibility while distinguishing near-term, implementable capabilities from advanced research functionalities, is depicted in Table 1.

Table 1. Staged deployment. CNN, NLG and XAI modules are indispensable to the architecture of the core system; federated learning, generative AI, and blockchain are introduced after proof-of-concept validation.

Stage	Scope	Core Components	Objective
Stage I (Pilot Implementation)	Single or small-network deployment within a controlled diagnostic laboratory environment	CNN + NLG + XAI (Gen-AI for initial model training)	Evaluate workflow integration, diagnostic accuracy and radiologist satisfaction, using WB-MRI for multiple myeloma
Stage II (Scaling)	Multi-institution collaboration	Federated Learning + expanded datasets	Develop privacy-preserving model generalization and cross-center standardization
Stage III (Future Expansion)	Research and regulatory optimization	Indication expansion Blockchain	Data augmentation, auditability and long-term data traceability

Whole Body MRI (WB-MRI)

The deployment of WB-MRI as the selected image-acquiring modality for the proposed framework is inextricably connected to a set of particular technical and diagnostic characteristics that differentiate it from traditional medical imaging techniques. It provides a detailed, non-invasive, ionizing radiation-free visualization of anatomical structures, in a single examination [20,21]. It offers adequate soft-tissue and bone marrow contrast, capable of depicting both spatial relationships and suggested disease activity, which is especially useful for early detection of neoplastic lesions or ongoing inflammatory processes [4]. WB-MRI protocols generate quantitative, reproducible signals suitable for AI-driven scoring, that follows predetermined established reporting systems, such as MY-RADS or MET-RADS-P. The high signal-to-noise and multiparametric image stacks allow for AI-compatible data substrates that enable smoother integration of the deep learning mechanisms into tasks such as denoising, semantic segmentation and lesion detection [8,21]. Furthermore, the procedural workflow is inherently suited to asynchronous interpretation, not requiring real-time radiologist assessment.

Nevertheless, WB-MRI still faces a number of challenges, including prolonged scan time (typically 30-50 minutes), which may increase patient discomfort or lead to motion artifacts (breathing, bowel peristalsis); occasionally inconsistent achievable homogeneity of the static and radiofrequency magnetic fields, resulting in distortions of the acquired images [20]; and increased operational costs, discouraging both institutional adoption and selection by patients.

Despite its limitations, WB-MRI often surpasses computed tomography, bone scintigraphy, and positron emission tomography in detecting and characterizing tumoral lesions [4,22]. As a consequence, international guidelines are gradually incorporating an increasing number of pathological entities that should or could be managed applying evolving WB-MRI techniques. Besides multiple myeloma, prostate cancer and hereditary cancer syndromes, there are suggestions of using WB-MRI as a valid imaging alternative in the staging and monitoring of melanoma, breast cancer, ovarian cancer, lymphoma, lung and colorectal cancer [7,8]. However, its application as an oncological screening tool in asymptomatic individuals remains debatable, albeit there is significant interest in the research community [23,24]. Furthermore, current imaging protocols cannot support routine screening in the general population for non-malignant conditions, including traumatic lesions, systemic inflammatory, musculoskeletal or metabolic disorders, or even surveillance imaging; the high false positive rate, including incidental findings leading to unnecessary interventions, along with the excessive cost of the exam are major constraints to the widespread adoption of this modality [23]. Yet, as improvements in scanner gradient systems, accelerated imaging, and sequence planning progress, WB-MRI could become feasible in a broader clinical setting.

For the time being, the CNN of the described conceptual framework can leverage currently available whole-body imaging technologies. Diffusion-weighted imaging (DWI) is indispensable to modern WB-MRI. Higher b-values allow for detailed apparent diffusion coefficient (ADC) maps, which enable the distinction between hypercellular tumors, dense marrow infiltration or abscesses, and necrotic lesions, cysts, oedema or areas of reduced cellularity, in a single-shot sequence. Disease assessment, metastasis detection and evaluation of response to treatment may be performed using either 1.5-tesla or 3-tesla scanners; the latter provide higher signal-to-noise reduction (SNR) but may be more prone to distortion artifacts, false positives, and patient overheating due to higher energy deposition [4]. T2-weighted images could be acquired using a fast spin-echo (FSE) technique to clearly highlight fluid-containing structures or, they could be further enhanced with a STIR (Short Tau Inversion Recovery) sequence —a T2-sensitive sequence with fat suppression—particularly useful for imaging the whole spine, vertebral metastases, or spinal cord compression. As an adjunctive sequence, T1-weighted images should preferably be obtained with gradient-echo (GRE) Dixon acquisitions, which tend to give more uniform fat suppression, improving visualization of anatomical detail. The axial plane should serve as the primary dataset, while coronal and sagittal reconstructions could enhance spatial context and improve lesion detection consistency.[4,25,26]

AI tools and technologies

Convolutional Neural Network (CNN)

The core AI component of the proposed platform, a CNN, is ideally suited for automated WB-MRI analysis, as it can evaluate hierarchical spatial features directly from volumetric data, can capture complex anatomical structures without the need for manual feature engineering, and often achieve stronger performance in classification and segmentation tasks compared to traditional machine learning approaches, like radiomics, that rely on handcrafted features [27]. Through the application of convolutional filters over localized receptive fields, CNNs can effectively identify variations in tissue contrast, and pathological patterns, essential for MRI slices analysis. This data-driven signal processing enables robust examination of non-linear locational relationships and morphological deformations, typical in medical imaging [28]. Table 2 presents examples of CNN-based models evaluated in MRI diagnostic tasks.

In medical imaging studies, CNNs have been applied to a broad spectrum of tasks, including detection of tumors, fractures or anatomical lesions (presence or absence) [9,29]. Especially when trained on sufficiently large and diverse cohorts, CNN-based systems have demonstrated significant diagnostic performance, particularly in brain tumor characterization, as well as multiple sclerosis and musculoskeletal assessment [12,29,30]. Furthermore, to effectively address challenges such as patient heterogeneity, and limited availability of annotated WB-MRI datasets [31,32], contemporary architectures incorporate multi-scale feature extraction, residual learning strategies, and encoder–decoder frameworks [10,28]. Within the dynamic concept of the proposed platform, the anticipated expansion in labeled data availability can gradually improve detection sensitivity, and enhance predictive reliability.

Table 2. CNN models already employed and evaluated in MRI diagnostic tasks

Examples of CNNs application in MRI scans interpretation and classification	
LeNet	Sarraf S. et al. [33]
AlexNet	Kumar L. S. et al. [34]
VGGNet (VGG16, VGG19)	Jain M. et al. [13]
GoogLeNet / Inception	Swarup C. et al. [35]
ResNet (e.g. ResNet-50)	Sun H. et al. [14]
DenseNet	Gottapu R. D. et al. [36]
U-Net and its variants (e.g., 3D U-Net, Attention U-Net)	Du G. et al. [15]

5.2. Natural Language Generation (NLG)

In the context of the automated analysis performed at the local radiological units, NLG can serve as the intermediary between the computational output of the CNN and the radiologist. Although convolutional algorithms can accurately identify and delineate suspicious lesions or pathological regions, their outputs are limited to segmentation masks representing localization of abnormal anatomical structures, probability scores, and quantitative measurements. Through NLG, these data can be formulated as meaningful, narrative text, describing the location, extent, and characteristics of imaging findings in a structured and clinically interpretable manner. This conversion supports faster assessment of multiple WB-MRI slices, and decreases the radiologist's workload, contributing to overall expert satisfaction. Additionally, the generation of standardized textual outputs allows for reporting homogeneity, reproducibility, integration within electronic health records, and interoperability among the collaborating health centers of the network [37,38].

From a technical point of view, NLG systems can either generate text from textual data ("text-to-text" generation), or generate text from non-linguistic data, such as medical images ("data-to-text" generation) – a viable approach for MRI slices interpretation. Specifically, advanced transformer-based large language models (LLMs) have demonstrated superior performance in documenting structured medical reports compared to previous generations, such as Long-Short Term Memory (LSTM) models. The nature of WB-MRI data can be effectively captured in vision-language LLM systems, where a vision encoder transforms image inputs into vector embeddings, while a language encoder processes textual inputs, and a decoder translates these representations into coherent human language [39].

Federated learning (FL)

The proposed platform envisions a polycentric FL architecture for AI-assisted analysis of WB-MRI images across multiple healthcare institutions. The nature of the design mandates that the contributing healthcare centers be equipped with state-of-the-art WB-MRI scanners, desirably open or semi-open to improve comfort of the examinee, reducing claustrophobia [40,41]. In a real-world setting, and in order to minimize heterogeneity of raw data due to geographic disparity and differences in scanner specifications, a suitable candidate to test the first operational use case could be a single established private network of diagnostic laboratories that would agree to include pilot use of the proposed platform in its daily workload of MRI exams.

After initial training a central CNN model is distributed to the participating institutions. The produced WB-MRI data is assessed with hybrid AI-augmented and human aforementioned workflow, and local model parameter updates are computed. Only the updates are transmitted to the federation server. The server aggregates updates from all institutions are aggregated. Finaly the improved model is redistributed.

Conceptually:

$$W_{global}^{t+1} = \sum_{k=1}^{K} \frac{n_k}{N} W_k^t$$

This represents the classical Federated Averaging (FedAvg) procedure [42].

Where:

W_k^t= local model from institution k
n_k= local dataset size
N= total number of samples
W_{global}^{t+1}= updated global model

Technically, model updates are transferred as tensors, i.e., multidimensional numerical arrays, and serialized binary parameter files. Common serialization formats include PyTorch (.pt), TensorFlow checkpoints, NumPy arrays, and protobuf-based communication packets [40-43]. Consequently, the exchanged information is numerical rather than visual or textual.

Anonymization - Encryption

Anonymization and encryption constitute valuable complementary mechanisms for patient privacy safeguarding and data security ensuring. [40,44]. Anonymization seeks to remove or obfuscate personally identifiable information (PII) from original data and associated metadata, thereby reducing the risk of patient re-identification during storage, sharing, and model development. Regarding WB-MRI, this process includes de-identification of DICOM headers, erasure of direct identifiers including patient's name, and elimination of indirect identifiers such as acquisition dates or institution-specific tags (e.g. the logo) [44,45]. Furthermore, advanced anonymization strategies address the threat of identity disclosure arising from the nature of imaging itself. For example, defacing or skull-stripping procedures may be applied to brain MRI scans to remove facial characteristics [44,46].

Encryption complements anonymization by protecting local data against unauthorized access throughout its lifecycle, including processing, transmission and storage. Cryptographic methods—such as symmetric and asymmetric encryption—are widely employed to protect data both at rest and in transit, thus ensuring confidentiality in distributed, federated medical imaging facilities [44,46]. Within the proposed diagnostic workflow, anonymization and encryption can prove to be particularly critical when WB-MRI datasets are transmitted between imaging devices like scanners, computational analysis systems and hospital information repositories, as such interfaces represent potential sites of vulnerability.

Explainable AI (XAI)

Although the combination of automated segmentation with natural language generation models may achieve high diagnostic performance in a structured meaningful report, their inherently opaque or "black box" characteristics can limit clinical confidence, and hinder integration into routine healthcare practice. To address this limitation, XAI methodologies can enable visualization and interpretation of model decision-making processes. Specifically, XAI approaches highlight image regions and salient features that contribute to diagnostic outputs, thereby providing insight into the rationale underlying model predictions [47]. In the WB-MRI examination paradigm, this application could be implemented through flagged MRI slices, highlighting the regions of interest that informed the prediction. This clear overview allows the radiologist to evaluate whether predictions are based on clinically meaningful imaging patterns, facilitating model validation, identification of systematic errors, such as artifacts, recognition of false positive findings, and alignment with established radiological knowledge [48].

As a consequence, XAI contributes to model optimization, by providing standardized, systematic and constructive performance assessment, across multiple, heterogeneous imaging protocols and diverse patient populations. This approach promotes effective human–machine collaboration, enabling active justification of algorithmic outputs rather than reliance on automated predictions [47,49]. From both clinical and regulatory perspectives, explainability is important for radiologist diagnostic confidence and accountability.

Generative Artificial Intelligence (GenAI)

The relative scarcity of annotated WB-MRI datasets, combined with restrictions on direct access to primary, officially diagnosed MRI scans, may position GenAI as a functional solution in the designed setting [50]. It could offer adjunctive support towards an adequate, satisfactory level of diagnostic accuracy, with methodologies including generative adversarial networks (GANs), variational autoencoders (VAEs), diffusion models (DMs), and transformer-based architectures [19,50]. These algorithms can generate novel samples, that approximate the underlying data distribution of medical images, thereby enabling counterfactual representations of anatomical lesions, and mimicking alternative pathological patterns. GANs, in particular, can strengthen CNN training and validation by harnessing probabilistic modeling and data synthesis. Once normal anatomical representations have been learned, subtle or rare abnormalities can be detected by identifying deviations from expected structural patterns [19,51].

Blockchain

The integration of blockchain technology into AI augmented diagnostic systems in routine clinical practice, although emerging, has not yet been well established. It has already been applied to support the secure storage and management of medical data through distributed cryptographic databases [52]. We, therefore, propose its consideration as a potential infrastructural component of the suggested network, for specific use cases. In instances involving unconfirmed diagnostic outputs or delayed identification of misdiagnosis, it becomes critical to implement mechanisms capable of tracing and verifying the provenance of these interactions. In such contexts, blockchain technology can provide data tracking and management, auditability, and trust-enhancing mechanisms [43]. It can also support longitudinal regulatory compliance monitoring by creating immutable or append-only records that document data access events, and record model updates or diagnostic predictions. The long term ability to trace the origin of a specific decision or action may prove to be critical for the appropriate attribution of responsibility or credit. Such functionalities may be particularly relevant in multi-parametric diagnostic ecosystems involving numerous institutions or stakeholders [52,53].

Expert-in-the-loop

Within the conceptualized framework, the radiologist holds a pivotal dual role; they are solely accountable for the finalized official diagnostic report that reaches the patient, and they contribute to the ongoing refinement of the deployed algorithms. The initial output of the local AI system consists of the raw WB-MRI slices and a set of tagged image slices delineating regions of interest (ROIs), before generating a draft diagnostic assessment. The resulting flagged lesion map is reviewed by the radiologist (detection-first approach), who interprets the findings, and decides whether to (a) confirm the lesion, (b) dismiss the target area as false positive or artefact, or (c) label it as uncertain. If uncertainty remains, a secondary consultation with another experienced radiologist, or even a multidisciplinary discussion is recommended; such events can be logged for subsequent model retraining. In all scenarios, the expert in charge ensures that the conclusive diagnosis aligns with clinical context and medical history, delivers the finalized report, and makes informed clinical suggestions, guiding management of the patient.

In this light, the algorithm enables automation of mundane, clerical tasks, such as image segmentation, reduction of noise-related artifacts and image reconstruction, volumetric and dimensional analysis, or standardized disease classification. Furthermore, it can detect minute morphological abnormalities and subtle image patterns that might, otherwise, remain undetected among the volume of WB-MRI images [54]. In other words, it acts as an augmented observer, assisting the radiologist with labor-intensive manual processes.

This can allow for fewer missed cases of disease —that is lower rate of false negatives— leading to improved diagnostic sensitivity, which can be used to evaluate success of the system. An expected reduction in turnaround time, defined as the interval between input of the original slices and delivery of the actual report to the patient is the second criterion of successful implementation. Neither of them, however, can guarantee satisfaction of the radiologist, which is of paramount importance for continuing utilization of the system, and is considered a third objective, that can be measured with post-implementation surveys.

Following delivery of the report to the patient, an equally important parameter of the proposed pipeline is image curation, labeling and annotation by the expert, who ensures that the source material which will be used for model training is representative and clinically meaningful. This data stewardship can be particularly useful, not only for preserving quality of the training process among the collaborating units of the infrastructure, but also for producing datasets appropriate for regulatory sandboxes [54], that can facilitate large scale multicenter national or international studies.

Key sociotechnical dimensions

Legal considerations

The foundation of a universal diagnostic network raises critical concerns of legitimacy, encompassing concerns of legal, ethical, and societal nature. The establishment of statutory and industry self-regulatory frameworks attempts to address issues regarding liability in AI implementation, post-deployment supervision, and dynamic algorithm updates. While deep learning models have demonstrated high performance in medical imaging interpretation, evidence shows that reproducibility of outcomes, critical protection of personal data, and transparent reporting of model limitations cannot be assumed [55]. Such conditions necessitate multilevel performance assessment, including internal and external validation against technical and diagnostic reference standards, and well-defined consent procedures. The delineation of responsibility among model developers, radiological institutions, and healthcare professionals remains, therefore, a complex and sensitive prerequisite.

In this regard, the European Union Artificial Intelligence Act [56], as the first-ever comprehensive legal framework on AI development and deployment worldwide, currently classifies AI-driven medical devices and software as high-risk AI systems, because they can influence decision-making and potentially cause harm. This designation means that a cascade of obligations for the developer, the provider and, to some extent, the deployer, should be fulfilled, in areas including risk management and data governance, transparency and human oversight, post-market surveillance, and audit-ready logging [54]. Within these lines, the proposed network functions as an AI-assisted decision-support tool, not an autonomous medical device; consequently, the radiologist retains final authority over clinical interpretations, aligning the system with risk-management obligations. The successful integration of innovative digital frameworks as diagnostic tools, like the suggested platform, depends not only on their effectiveness, but also largely on incorporating safety mechanisms to ensure compliance with the applicable legislation and, most importantly, to warrant their accountable and equitable application in clinical practice.

Fairness

Fairness in AI refers to the non-discriminatory treatment of all patient populations, ensuring that model performance remains consistent across diverse demographic patient groups, including variations in age, gender, ethnicity, and socioeconomic status [57]. Algorithmic bias may arise from unrepresentative training datasets or when the underlying data distribution alters the input space. It may also occur due to flawed labeling processes, or unbalanced model updating and training. When this performance drift occurs unevenly across demographic subgroups or clinical contexts, it becomes a source of discriminatory bias, even if the model was originally well-calibrated. Such biases can lead to diagnostic inaccuracies, disproportionately affecting marginalized populations and perpetuating healthcare disparities [55,57,58]. Within this context, we hypothesize that the FL architecture of the proposed framework will support a satisfactory degree of fair and equal performance, particularly when deployed in conjunction with rigorous dataset curation, systematic cross-demographic performance auditing, and transparent retraining protocols.

AI-related anxiety

The application of comprehensive whole-body medical examinations, which carry the potential of revealing previously undetected and possibly life-threatening conditions, introduces important psychological dimensions that may substantially affect patient well-being. Additionally, certain procedural aspects of the MRI examination itself, such as prolonged duration, noisy or confined environments, often become a source of emotional distress. The integration of AI-enhanced imaging systems that limit direct interaction with the healthcare professional may further contribute to uncertainty, leading to anxiety, especially when outputs are perceived as opaque or insufficiently explained. Evidence derived from studies employing validated psychometric parameters [59],[60] indicates that anticipatory anxiety may increase when medical imaging interpretations are delivered or mediated through electronic accessories, without adequate contextualization by the healthcare expert [61]. In the case of the proposed solution, these effects can be mitigated when the patient is well informed on the process-specific features of WB-MRI, and the benefits of this particular imaging modality within the context of their specific pathology.

In pilot implementation studies, anxiety associated with AI applications has been investigated alongside technical performance metrics using structured, longitudinal designs. Standardized instruments, such as the State-Trait Anxiety Inventory (STAI) and the Hospital Anxiety and Depression Scale (HADS), can be administered at predefined time points (e.g., pre-exposure, post-diagnosis, and follow-up) to quantify transient and enduring psychological responses to the workflow of the proposed platform [62]. Within the platform operation, these outcomes may be further evaluated using questionnaires that assess perceived transparency, diagnostic reliability, and trust. Structured patient–clinician communication protocols could be incorporated as predefined intervention components. Consequently, additional endpoints might include changes in anxiety scores, while also evaluating compliance to treatment and overall patient satisfaction. Such provisions may enable systematic evaluation of psychosocial outcomes, supporting the development of evidence-based operational guidelines.

DISCUSSION

The presented framework demonstrates how phased implementation can reconcile technological aspiration with clinical feasibility. An initial, mono-institutional pilot using WB-MRI for multiple myeloma could offer measurable endpoints and manageable data governance; once effectiveness and usability are demonstrated, federated expansion across multiple centers can proceed. Yet, if "data is the new oil" as Clive Humby pointed out in the early 2000s [63,64], then data repositories are the foundational reservoirs of algorithmic "fuel" for any AI-assisted medical imaging application, such as the one proposed. Establishing comprehensive image-and-text-centric digital ecosystems, across national or international entities, with large quantities of properly labeled medical images can solve the key question of data fragmentation. Safe and efficient testing of radiology related algorithms, under transparent and rigorous evaluation, requires leveraging structured, large-scale, and high-quality source material, in regulated frameworks that safeguard fundamental rights, without stifling innovation. WB-MRI datasets, which are necessary for supervised training and validation of the selected CNN, are considered of especially "high-value", due to their length and complexity. While the suggested setup may initially rely on suitable datasets deposited by the participating healthcare units, it will greatly benefit from such curated environments, that can vastly improve its predictive capacity.

Following the stipulation of EU AI Act for geographically and contextually relevant datasets [56], the European Health Data Space (EHDS) constitutes a groundbreaking framework that will enable the composition of GDPR (General Data Protection Regulation)-compliant datasets, shared across 27 member-states. Under its auspices, the EUropean Federation for CAncer IMages (EUCAM) initiative serves as the primary European legal hub for WB-MRI data, included in a repository that hosts more than 100,000 radiology cases and 60 million medical images; these scans come with clinical "ground truth" (biopsies, follow-ups, and radiologist labels) [65]. Other sources can be found in the UK Biobank [17], the German National Cohort (NAKO) [18], and The Cancer Imaging Archive (TCIA) [66], all of which hold large archives of standardized WB-MRI images in DICOM format [67]. Availability readiness of the data varies greatly from Open Access in TCIA, to obligation of training only within a specific secure cloud workspace (Trusted Research Environment) in EUCAIM. Data governance remains a key parameter influencing the success of the proposed system.

One of the most obvious failure scenarios of the network, caused by inadequate access to useful raw material, involves demographic and spectrum bias. If the model's training recycles a narrow range of disease cases or population subsets, it may show higher error rates in the actual clinical setting, where individual patients may vary considerably. Over time, and especially in view of the everchanging population landscape, this may degrade its performance, undermining its reliability and eroding the trust of the human constituent. The misalignment in human-AI interaction represents another possible failure scenario. For instance, alert fatigue resulting from too many false positives may lead to algorithmic aversion, potentially missing a true positive when it actually occurs. Moreover, it has been shown that AI predictions with large errors tend to lead to negative treatment effects, suggesting that when the radiologist struggles to consistently distinguish between accurate and inaccurate predictive AI outputs, ultimately the accuracy of the synergistic system diminishes [68]. Furthermore, the performance characteristics of the AI system itself represent a major factor that may render the system prone to failure; lower AI confidence levels —that is uncertainty in the provided prompts— or poorly generated radiology reports, lacking required nuance, are bound to increase turnaround times and, at the same time, increase the radiologist mental workload.

The complex and often unpredictable dynamics between radiologists and AI-driven applications could guide future research towards exploring the need for individualized approaches that are aware of clinician heterogeneity as well as the limitations of the adaptive algorithms [68]. Collaboration between radiologists and AI developers can allow for improved workflow integration, that strengthens rather than hinders the expert's confidence and effectiveness. Further investigation is required to determine optimal hybrid CNN-NLG architectures for WB-MRI image interpretation, and to define the appropriate extent of local model updating prior to parameter sharing within the FL framework. Large-scale longitudinal studies or even carefully laid experimental setups could also investigate the actual predictors of successful synergy between the human expert and the AI component, playing an instrumental role in the maturity of diagnostic platforms like the one proposed.

More readily manageable in the foreseeable future is the development of training programs that will enable the radiologist to properly understand the AI systems being deployed, so as to correctly interpret and contextualize their output. In light of the complexities and challenges surrounding the integration of adaptive algorithms into radiology, AI literacy emerges as a clinical imperative in order to ensure adequate human oversight and retain expert's accountability. In the same setting, communicating clearly with the examinee is essential for maintaining trust and transparency in patient care. Nonetheless, the approach to informed patient consent presents itself as a complex issue that requires future consideration and dialogue [59].

CONCLUSION

The present proposition focuses on a human centered, AI-assisted diagnostic framework that integrates recent advances in WB-MRI techniques and a multimodal vision-language model into a seamless workflow, where continuous feedback between the radiologist and the adaptive algorithm enables the progressive refinement of model performance, without compromising clinical reliability. The clinical boundary of its deployment is defined by the specific indications of use, while its technical boundary is linked to the immaturity of the synthetic intelligence, limiting the level of automation. Such a platform is feasible but conditional upon the interplay of technological innovation and data accessibility, legal and ethical governance. There is no doubt that delegating critical diagnostic judgments to automated systems without clearly establishing liability not only raises legitimacy concerns but may also lead to gradual erosion of human accountability. The implementation of the hypothesized framework could serve as a compelling paradigm of human-AI synergy, where computational precision and clinical intuition converge to a sustainable solution, with ultimate objective to optimize the healthcare quality provided to the individual patient.

Author Contributions: Conceptualization, investigation, methodology, writing, C.Z.; project administration, validation, H.L.; supervision, H.L., C.P. and G.G.; All authors have read and agreed to the published version of the manuscript.

Funding: The work presented in this article was funded by the special account for research grants of the University of West Attica.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Conflicts of Interest: The authors have no conflicts of interest to declare that are relevant to the content of this article.

REFERENCES

1. Kelly, B.S.; Judge, C.; Bollard, S.M.; et al. Radiology artificial intelligence: A systematic review and evaluation of methods (RAISE). Eur. Radiol. 2022, 32, 7998–8007. https://doi.org/10.1007/s00330-022-08784-6.
2. Alami, H.; Lehoux, P.; Auclair, Y.; de Guise, M.; Gagnon, M.P.; Shaw, J.; Roy, D.; Fleet, R.; Ag Ahmed, M.A.; Fortin, J.P. Artificial intelligence and health technology assessment: Anticipating a new level of complexity. J. Med. Internet Res. 2020, 22, e17707. https://doi.org/10.2196/17707.
3. Petticrew, M.; Sowden, A.; Lister-Sharp, D. False-negative results in screening programs: medical, psychological, and other implications. Int. J. Technol. Assess. Health Care 2001, 17, 164–170. https://doi.org/10.1017/S0266462300105021
4. Summers, P.; Saia, G.; Colombo, A.; Pricolo, P.; Zugni, F.; Alessi, S.; Marvaso, G.; Jereczek-Fossa, B.A.; Bellomi, M.; Petralia, G. Whole-body magnetic resonance imaging: Technique, guidelines and key applications. Ecancermedicalscience 2021, 15, 1164. https://doi.org/10.3332/ecancer.2021.1164.
5. Fritz, B.; Fritz, J. Artificial intelligence for MRI diagnosis of joints: A scoping review of the current state-of-the-art of deep learning-based approaches. Skeletal Radiol. 2022, 51, 315–329. https://doi.org/10.1007/s00256-021-03830-8.
6. Nawaz, A.; Rehman, A.U.; Ali, T.M.; Hayat, Z.; Rahim, A.; Uz Zaman, U.K.; Ali, A.R. A comprehensive literature review of application of artificial intelligence in functional magnetic resonance imaging for disease diagnosis. Appl. Artif. Intell. 2021, 35, 1420–1438. https://doi.org/10.1080/08839514.2021.1982185.
7. Morone, M.; Bali, M.A.; Tunariu, N.; Messiou, C.; Blackledge, M.; Grazioli, L.; Koh, D.M. Whole-body MRI: Current applications in oncology. AJR Am. J. Roentgenol. 2017, 209, W336–W349. https://doi.org/10.2214/AJR.17.17984.
8. Zugni, F.; Mariani, L.; Lambregts, D.M.J.; et al. Whole-body MRI in oncology: Acquisition protocols, current guidelines, and beyond. Radiol. Med. 2024, 129, 1352–1368. https://doi.org/10.1007/s11547-024-01851-6.
9. Shimron, E.; Perlman, O. AI in MRI: Computational Frameworks for a Faster, Optimized, and Automated Imaging Workflow. Bioengineering 2023, 10, 492. https://doi.org/10.3390/bioengineering10040492
10. Wang, T.W.; Hong, J.S.; Lee, W.K.; et al. Performance of convolutional neural network models in meningioma segmentation in magnetic resonance imaging: A systematic review and meta-analysis. Neuroinformatics 2025, 23, 14. https://doi.org/10.1007/s12021-024-09704-3.

11. Gheibi, Y.; Shirini, K.; Razavi, S.N.; et al. CNN-Res: Deep learning framework for segmentation of acute ischemic stroke lesions on multimodal MRI images. BMC Med. Inform. Decis. Mak. 2023, 23, 192. https://doi.org/10.1186/s12911-023-02289-y.
12. Luo, C.; Li, X.; Wang, L.; He, J.; Li, D.; Zhou, J. How Does the Dataset Affect CNN-Based Image Classification Performance? Proceedings of the 2018 5th International Conference on Systems and Informatics (ICSAI), Nanjing, China, 10–12 November 2018; pp. 361–366. https://doi.org/10.1109/ICSAI.2018.8599448 (accessed on April 23, 2026)
13. Jain, M.; Kaushik, P.; Shah, A. Comparison of VGG16 and VGG19 convolutional neural network (CNN) layers on MRI brain tumor detection. Int. J. Trend Sci. Res. Dev. 2016, 1, 275–280. https://www.ijtsrd.com/papers/ijtsrd3542.pdf (accessed on April 4, 2026)
14. Sun, H.; Wang, A.; Wang, W.; Liu, C. An improved deep residual network prediction model for the early diagnosis of Alzheimer's disease. Sensors 2021, 21, 4182. https://doi.org/10.3390/s21124182
15. Du, G.; Cao, X.; Liang, J.; Chen, X.; Zhan, Y. Medical image segmentation based on U-net: A review. J. Imaging Sci. Technol. 2020, 64, 020508. https://doi.org/10.2352/J.ImagingSci.Technol.2020.64.2.020508
16. Sudlow, C.; Gallacher, J.; Allen, N.; Beral, V.; Burton, P.; Danesh, J.; Downey, P.; Elliott, P.; Green, J.; Landray, M.; Liu, B.; Matthews, P.; Ong, G.; Pell, J.; Silman, A.; Young, A.; Sprosen, T.; Peakman, T.; Collins, R. UK Biobank: An Open Access Resource for Identifying the Causes of a Wide Range of Complex Diseases of Middle and Old Age. PLoS Med. 2015, 12, e1001779. https://doi.org/10.1371/journal.pmed.1001779
17. UK Biobank. Health Research Data for the World. Available online: https://www.ukbiobank.ac.uk/ (accessed on April 27, 2026)
18. German National Cohort (NAKO). German National Cohort: Germany's Largest Long-Term Study to Research Widespread Diseases. Available online: https://nako.de/en/ (accessed on April 25, 2026)
19. Ibrahim, M.; Khalil, Y.A.; Amirrajab, S.; Sun, C.; Breeuwer, M.; Pluim, J.; Elen, B.; Ertaylan, G.; Dumontier, M. Generative AI for synthetic data across multiple medical modalities: A systematic review of recent developments and challenges. Comput. Biol. Med. 2025, 189, 109834. https://doi.org/10.1016/j.compbiomed.2025.109834
20. Kwee, R.M.; Kwee, T.C. Whole-body MRI for preventive health screening: A systematic review of the literature. J. Magn. Reson. Imaging 2019, 50, 1489–1503. https://doi.org/10.1002/jmri.26736
21. Cosmus, T.C.; Parizh, M. Advances in whole-body MRI magnets. IEEE Trans. Appl. Supercond. 2011, 21, 2104–2109. https://doi.org/10.1109/TASC.2010.2084981
22. Tunariu, N.; Blackledge, M.; Messiou, C.; Petralia, G.; Padhani, A.; Curcean, S.; Curcean, A.; Koh, D.M. What's new for clinical whole-body MRI (WB-MRI) in the 21st century. Br. J. Radiol. 2020, 93, 20200562. https://doi.org/10.1259/bjr.20200562
23. Dai, K.Z.; Jambawalikar, S.R.; Kang, S.K. Whole-Body MRI Screening of Average Risk Populations: Promises and Controversies. J. Magn. Reson. Imaging 2026, 63, 1217–1223. https://doi.org/10.1002/jmri.70268
24. Kierans, A.S.; Hentel, K.D.; Dodelzon, K.; Shih, G.; Prince, M.R.; Lantos, J.; Frey, M.K.; Sharaf, R.N.; Min, R.J.; Guniganti, P. How to Implement a Radiologist Led Whole-Body MRI Screening Program. Radiol. Adv. 2026, 3, umag014. https://doi.org/10.1093/radadv/umag014
25. Ozkara, B.B.; Chen, M.M.; Federau, C.; Karabacak, M.; Briere, T.M.; Li, J.; Wintermark, M. Deep Learning for Detecting Brain Metastases on MRI: A Systematic Review and Meta-Analysis. Cancers 2023, 15, 334. https://doi.org/10.3390/cancers15020334
26. Grøvik, E.; Yi, D.; Iv, M.; Tong, E.; Rubin, D.; Zaharchuk, G. Deep Learning Enables Automatic Detection and Segmentation of Brain Metastases on Multisequence MRI. J. Magn. Reson. Imaging 2020, 51, 175–182. https://doi.org/10.1002/jmri.26766
27. Ottoni, M.; Kasperczuk, A.; Tavora, L.M.N. Machine Learning in MRI Brain Imaging: A Review of Methods, Challenges, and Future Directions. Diagnostics 2025, 15, 2692. https://doi.org/10.3390/diagnostics15212692
28. Litjens, G.; Kooi, T.; Ehteshami Bejnordi, B.; et al. A survey on deep learning in medical image analysis. Med. Image Anal. 2017, 42, 60–88. https://doi.org/10.1016/j.media.2017.07.005
29. Zachariadis, C.B.; Leligou, H.C. Harnessing artificial intelligence for automated diagnosis. Information 2024, 15, 311. https://doi.org/10.3390/info15060311.
30. Cho, J.; Lee, K.; Shin, E.; Choy, G.; Do, S. How much data is needed to train a medical image deep learning system to achieve necessary high accuracy? arXiv 2015, arXiv:1511.06348. https://doi.org/10.48550/arXiv.1511.06348
31. Luu, M.S.K.; Benedichuk, M.V.; Roppert, E.I.; Kenzhin, R.M.; Tuchinov, B.N. A structured review and quantitative profiling of public brain MRI datasets for foundation model development. J. Imaging 2025, 11, 454. https://doi.org/10.3390/jimaging11120454

32. Dishner, K.A.; McRae-Posani, B.; Bhowmik, A.; Jochelson, M.S.; Holodny, A.; Pinker, K.; Eskreis-Winkler, S.; Stember, J.N. A survey of publicly available MRI datasets for potential use in artificial intelligence research. J. Magn. Reson. Imaging 2024, 59, 450–480. https://doi.org/10.1002/jmri.29101
33. Sarraf, S.; Tofighi, G. Classification of Alzheimer's disease structural MRI data by deep learning convolutional neural networks. arXiv 2016, arXiv:1607.06583. https://doi.org/10.48550/arXiv.1607.06583
34. Kumar, L.S.; Hariharasitaraman, S.; Narayanasamy, K.; Thinakaran, K.; Mahalakshmi, J.; Pandimurugan, V. AlexNet approach for early stage Alzheimer's disease detection from MRI brain images. Mater. Today Proc. 2022, 51, 58–65. https://doi.org/10.1016/j.matpr.2021.04.415.
35. Swarup, C.; Singh, K.U.; Kumar, A.; Pandey, S.K.; Varshney, N.; Singh, T. Brain tumor detection using CNN, AlexNet and GoogLeNet ensembling learning approaches. Electron. Res. Arch. 2023, 31, 2900–2924. https://doi.org/10.3934/era.2023146
36. Gottapu, R.D.; Dagli, C.H. DenseNet for anatomical brain segmentation. Procedia Comput. Sci. 2018, 140, 179–185. https://doi.org/10.1016/j.procs.2018.10.327
37. Heo, T.S.; Kim, Y.S.; Choi, J.M.; Jeong, Y.S.; Seo, S.Y.; Lee, J.H.; Jeon, J.P.; Kim, C. Prediction of Stroke Outcome Using Natural Language Processing-Based Machine Learning of Radiology Report of Brain MRI. J. Pers. Med. 2020, 10, 286. https://doi.org/10.3390/jpm10040286
38. Lyu, M.; Li, X.; Chen, Z.; Pan, J.; Peng, C.; Talankar, S.; Wu, Y. Natural Language Generation in Healthcare: A Review of Methods and Applications. J. Biomed. Inform. 2026, 176, 104997. https://doi.org/10.1016/j.jbi.2026.104997
39. Pan, S.; Luo, L.; Wang, Y.; Chen, C.; Wang, J.; Wu, X. Unifying Large Language Models and Knowledge Graphs: A Roadmap. IEEE Transactions on Knowledge and Data Engineering 2024, 36(7), 3580–3599. https://doi.org/10.1109/TKDE.2024.3352100
40. Kaissis, G.A.; Makowski, M.R.; Rückert, D.; et al. Secure, privacy-preserving and federated machine learning in medical imaging. Nat. Mach. Intell. 2020, 2, 305–311. https://doi.org/10.1038/s42256-020-0186-1
41. Kwak, L.; Bai, H. The role of federated learning models in medical imaging. Radiol. Artif. Intell. 2023, 5, e230136. https://doi.org/10.1148/ryai.230136
42. Paszke, A.; Gross, S.; Massa, F.; et al. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In Advances in Neural Information Processing Systems 32 (NeurIPS 2019); 2019; pp. 8024–8035. https://doi.org/10.48550/arXiv.1912.01703
43. Passerat-Palmbach, J.; Farnan, T.; Miller, R.; Gross, M.S.; Flannery, H.L.; Gleim, B. A blockchain-orchestrated federated learning architecture for healthcare consortia. arXiv 2019, arXiv:1910.12603. https://doi.org/10.48550/arXiv.1910.12603
44. Zhu, Y.; Yin, X.; Wee-Chung Liew, A.; Tian, H. Privacy-preserving in medical image analysis: A review of methods and applications. In Parallel and Distributed Computing, Applications and Technologies (PDCAT 2024); Li, Y.; Zhang, Y.; Xu, J., Eds.; Lecture Notes in Computer Science, Vol. 15502; Springer: Singapore, 2025. https://doi.org/10.1007/978-981-96-4207-6_15
45. European Parliament and Council of the European Union. Regulation (EU) 2016/679 of the European Parliament and of the Council (General Data Protection Regulation). Off. J. Eur. Union 2016, L119, 1–88. https://eur-lex.europa.eu/legal-content/EL/TXT/?uri=celex:32016R0679 (accessed on April 25, 2026)
46. Diaz, O.; Kushibar, K.; Osuala, R.; Linardos, A.; Garrucho, L.; Igual, L.; Radeva, P.; Prior, F.; Gkontra, P.; Lekadir, K. Data preparation for artificial intelligence in medical imaging: A comprehensive guide to open-access platforms and tools. Phys. Med. 2021, 83, 25–37. https://doi.org/10.1016/j.ejmp.2021.02.007
47. Houssein, E.H.; Gamal, A.M.; Younis, E.M.G.; et al. Explainable artificial intelligence for medical imaging systems using deep learning: A comprehensive review. Cluster Comput. 2025, 28, 469. https://doi.org/10.1007/s10586-025-05281-5
48. Champendal, M.; Müller, H.; Prior, J.O.; Dos Reis, C.S. A scoping review of interpretability and explainability concerning artificial intelligence methods in medical imaging. Eur. J. Radiol. 2023, 169, 111159. https://doi.org/10.1016/j.ejrad.2023.111159
49. Hafeez, Y.; Memon, K.; Al-Quraishi, M.S.; Yahya, N.; Elferik, S.; Ali, S.S.A. Explainable AI in diagnostic radiology for neurological disorders: A systematic review, and what doctors think about it. Diagnostics 2025, 15, 168. https://doi.org/10.3390/diagnostics15020168
50. Alshanbari, A.H.; Alzahrani, S.M. Generative AI for diagnostic medical imaging: A review. Curr. Med. Imaging 2025, 21, e15734056369157. https://doi.org/10.2174/0115734056369157250212095252
51. Karmakar, A.; Shaw, A.; Rakshit, S.; et al. The role of generative AI in medical image synthesis: A review. Discov. Appl. Sci. 2025, 7, 1219. https://doi.org/10.1007/s42452-025-07714-7.
52. European Society of Radiology (ESR). ESR white paper: Blockchain and medical imaging. Insights Imaging 2021, 12, 82. https://doi.org/10.1186/s13244-021-01029-y.

53. Kasralikar, P.; Polu, O.R.; Chamarthi, B.; Rupavath, R.V.; Patel, S.; Tumati, R. Blockchain for securing AI-driven healthcare systems: A systematic review and future research perspectives. Cureus 2025, 17, e83136. https://doi.org/10.7759/cureus.83136
54. Kotter, E.; D'Antonoli, T.A.; Cuocolo, R.; et al. Guiding AI in Radiology: ESR's Recommendations for Effective Implementation of the European AI Act. Insights Imaging 2025, 16, 33. https://doi.org/10.1186/s13244-025-01905-x
55. Zhang, J.; Zhang, Z. Ethics and governance of trustworthy medical artificial intelligence. BMC Med. Inform. Decis. Mak. 2023, 23, 7. https://doi.org/10.1186/s12911-023-02103-9.
56. European Parliament; Council of the European Union. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence and Amending Regulations and Directives. Off. J. Eur. Union 2024, L, 2024/1689. Available online: https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=OJ:L_202401689 (accessed on April 25, 2026)
57. Ueda, D.; Kakinuma, T.; Fujita, S.; et al. Fairness of artificial intelligence in healthcare: Review and recommendations. Jpn. J. Radiol. 2023, 42, 3–15. https://doi.org/10.1007/s11604-023-01474-3.
58. Chinta, S.V.; Wang, Z.; Palikhe, A.; Zhang, X.; Kashif, A.; Smith, M.A.; et al. AI-driven healthcare: A review on ensuring fairness and mitigating bias. PLOS Digit. Health 2025, 4, e0000864. https://doi.org/10.1371/journal.pdig.0000864.
59. Masalma, R.; Zidan, T.; Amasheh, S.; Maree, M.; Alhanbali, M.; Shawahna, R. Predictors of Anxiety in Patients Undergoing Magnetic Resonance Imaging Scans: A Multicenter Cross-Sectional Study. BMC Psychiatry 2024, 24, 633. https://doi.org/10.1186/s12888-024-06091-6
60. Fakes, K. Patient Experiences and Anxiety Related to Medical Imaging: Challenges and Potential Solutions. J. Med. Radiat. Sci. 2024, 71, 3–6. https://doi.org/10.1002/jmrs.748
61. Yang, H.; Sundar, S.S. AI anxiety: Explication and exploration of effect on state anxiety when interacting with AI doctors. Comput. Hum. Behav. Artif. Hum. 2025, 3, 100128. https://doi.org/10.1016/j.chbah.2025.100128.
62. Teelhawod, B.N.; Akhtar, F.; Heyat, M.B.B.; Tripathi, P.; Mehrotra, R.; Asfaw, A.B.; Shorman, O.; Masadeh, M. Machine learning in e-health: A comprehensive survey of anxiety. In Proc. 2021 Int. Conf. Data Analytics for Business and Industry (ICDABI); Sakheer, Bahrain, 2021; pp. 167–172. https://doi.org/10.1109/ICDABI53623.2021.9655966.
63. Humby, C. Data is the new oil. In Proc. ANA Senior Marketer's Summit; Evanston, IL, USA, 2006.
64. Wikibooks Contributors. Lentis/"Data Is the New Oil".
Available online: https://en.wikibooks.org/wiki/Lentis/%22Data_is_the_new_oil%22 (accessed on April 26, 2026)
65. European Commission. European Cancer Imaging Initiative.
Available online: https://digital-strategy.ec.europa.eu/en/policies/cancer-imaging (accessed on April 25, 2026)
66. The Cancer Imaging Archive (TCIA). The Cancer Imaging Archive.
Available online: https://www.cancerimagingarchive.net/ (accessed on April 25, 2026)
67. Aiello, M.; Esposito, G.; Pagliari, G.; Borrelli, P.; Brancato, V.; Salvatore, M. How does DICOM support big data management? Investigating its use in medical imaging community. Insights Imaging 2021, 12, 164. https://doi.org/10.1186/s13244-021-01081-8
68. Yu, F.; Moehring, A.; Banerjee, O.; Salz, T.; Agarwal, N.; Rajpurkar, P. Heterogeneity and Predictors of the Effects of AI Assistance on Radiologists. Nat. Med. 2024, 30, 837–849. https://doi.org/10.1038/s41591-024-02850-w

